arXiv:1603.07751v1 [cs.GT] 24 Mar 2016

Equilibrium Selection in Information Elicitation without Verification via Information Monotonicity

Yuqing Kong
University of Michigan

Grant Schoenebeck
University of Michigan

Abstract

Peer-prediction is a mechanism which elicits privately-held, non-variable information from self-interested agents—formally, truth-telling is a strict Bayes Nash equilibrium of the mechanism. The original Peer-prediction mechanism suffers from two main limitations: (1) the mechanism must know the "common prior" of agents' signals; (2) additional undesirable and non-truthful equilibria exist which often have a greater expected payoff than the truth-telling equilibrium. A series of results has successfully weakened the known common prior assumption. However, the equilibrium multiplicity issue remains a challenge.

In this paper, we address the above two problems. In the setting where a common prior exists but is not known to the mechanism we show (1) a general negative result applying to a large class of mechanisms showing truth-telling can never pay strictly more in expectation than a particular set of equilibria where agents collude to "relabel" the signals and tell the truth after relabeling signals; (2) provide a mechanism that has no information about the common prior but where truth-telling pays as much in expectation as any relabeling equilibrium and pays strictly more than any other symmetric equilibrium; (3) moreover in our mechanism, if the number of agents is sufficiently large, truth-telling pays similarly to any equilibrium close to a "relabeling" equilibrium and pays strictly more than any equilibrium that is not close to a relabeling equilibrium.

1 Introduction

User feedback requests (e.g Ebay's reputation system and the innumerable survey requests in one's email inbox) are increasingly prominent and important. However, the overwhelming number of requests can lead to low participation rates, which in turn may yield unrepresentative samples. To encourage participation, a system can reward people for answering requests. But this may cause perverse incentives: some people may answer a large of number of questions simply for the reward and without making any attempt to answer accurately. In this case, the reviews the system obtains may be inaccurate and meaningless. Moreover, people may be motivated to lie when they face a potential loss of privacy or can benefit in the future by lying now.

It is thus important to develop systems that motivate honesty. If we can verify the information people provide in the future (e.g via prediction markets), we can motivate honesty via this future verification. However, sometimes we need to elicit information without verification since the objective truth is hard to access or even does not exist (e.g. a self-report survey for involvement in crime). In our paper, we focus on the situation where the objective truth is not observable.

One important framework for designing incentive systems without verification is *peer prediction* [12]. *Peer prediction* uses each person's information to predict other people's information and pays according to how good the prediction is. *Peer prediction* assumes people's information is related and the systems and the people share a common prior.

Despite this clever insight, information elicitation mechanisms, such a peer prediction, are not often used in practice. Two main obstacles that prevent peer prediction from being practical are that (1) the mechanism must know the prior; (2) there are many other equilibria besides truth-telling, which may pay as much or even more than the truth-telling equilibrium.

Recent research [6] indicates that individuals in lab experiments do not always truth-tell when faced with peer prediction mechanisms; this may in part be related to the issue of equilibrium multiplicity.

Our Contributions

1. We identify a particular set of equilibria, which we call *permutation equilibria*, which exist in every (non-trivial) information elicitation mechanism; and show that for any (non-trivial) information elicitation mechanism and for any permutation equilibrium, there exists a setting where that permutation equilibrium is paid at least as much as truth-telling;

2. We propose an information elicitation mechanism that we call *Disagreement Mechanism* where truth-telling is a Bayesian Nash equilibrium for any finite set of signals and any number of agents greater than 3, and where:

    a. Truth-telling is paid strictly more in expectation than any other symmetric equilibrium that is not a permutation equilibrium, which have an expected payoff equal to the truth-telling equilibrium;

    b. Any symmetric equilibrium that has expected payoff close to truth-telling must be "close" to a permutation equilibrium;

    c. If the number of agents is sufficiently large then any asymmetric equilibrium which is "close" to a permutation equilibrium has expected payoff "close" to that of the truth-telling equilibrium, and any equilibrium that is not "close" to a permuation equilibrium pays strictly less then the truth-telling equilibrium.

Permutation equilibria are intuitively unnatural and risky as they require extreme coordination amongst the agents. Thus our results about symmetric equilibrium are quite strong, despite the impossibility result. Asymmetric equilibrium require more coordination between agents than symmetric equilibrium. Additionally, the possible gains from doing so are limited and go to zero as the number of agents increase. We believe the mechanism we propose can motivate a small number of agents to tell the truth no matter how many signals they receive.

High Level Techniques Our *Disagreement Mechanism* pays agents individually (locally) for "agreement" and globally for "disagreement". When agents collude, they share information about their strategies. Since they are paid individually for "agreement", they will use their information about other people's strategies and "agree to agree" which reduces their global payoff which depends on "disagreement".

Essentially, when agents collude, our *Disagreement Mechanism* encourages each agent to implicitly admit their collusion by unilaterally increasing their individual payoff for doing so, but the mechanism then simultaneously decreases the total payoffs to all agents. Only when agents do not choose to collude and lack information about other people's strategies, can they "agree to disagree." In this case, even when they maximize their individual payoff, globally they still have a lot of disagreement.

Technical Contributions In addition to the above results, our works has several contributions in the techniques employed:

1. Our *Disagreement Mechanism* encourages not only agents with the same private information to agree, but also agents with different private information to disagree. We present a novel way to measure the amount of "information" by casting the reports as multiple labelled points in a space and measuring the quality of the "classification".

2. To show that the "classification" quality always decreases with non-truthful equilibria, we exploit tools from information theory, namely *Information Monotonicity*. Despite their natural and powerful application, to our knowledge, this is the first time such tools have been explicitly employed in the peer prediction literature.

3. We provide a framework to analyse the structure of the equilibrium of a very complex game.

1.1 Related Work

There are several papers [14, 18, 15, 16, 22, 17, 5, 19, 20, 21] that focused on the goal of removing the assumption that the mechanism knows the common prior and successfully weaken this known common prior assumption after peer prediction [12] is introduced.

Bayesian Truth Serum (BTS) [14] first successfully weakened the known common prior assumption and provides an important framework for mechanisms without known common prior. BTS requires the agents report—in addition to their reported signal—a forecast (prediction) of the other agents' reported signals, and uses this predictions in lieu of the common prior. BTS incentives agents to report accurate forecasts by rewarding forecasts that have the ability to predict the other agents' reported signal.

However, BTS has two weakness: (1) BTS requires that the number of agents goes to infinity (or is large enough in a modified version) since the mechanism needs agents to believe it has access to the true distribution of from which agents' signals are drawn. (2) The equilibrium analysis provided in [14] is in the case where the number of agents goes to infinity and only proves that truth-telling has expected payment at least as high as other equilibrium and so does not avoid the case that truth-telling pays equivalently to many other equilibrium.

Several mechanisms [18, 15, 16, 17, 19, 20, 21] are based on the BTS framework and address the first weakness of BTS. *Robust Bayesian Truth Serum (RBTS)* [18] is a mechanism which can only be applied to binary signals and the *multi-valued RBTS* [15] mechanism can be applied to non-binary signals while it requires an additional assumption that every agent believes the probability that other agents receive signal σ is higher if he himself also receives σ. Both of these works do not solve the equilibrium multiplicity issue, but do work for a small number of agents. *Minimal Truth Serum (MTS)* [17] is a mechanism where agents have the option to report or not report their predictions, and also lacks analysis of non-truthful equilibria. MTS uses a typical zero-sum technique such that all equilibria are paid equally. In contrast, we show that in our *Disagreement Mechanism* any equilibrium that is even close to paying more than the truth-telling equilibrium must be close to a small set of "permutation" equilibrium.

Thus, while the above work addresses the first weakness of BTS, it does not address the second and also lacks an analysis of the non-truthful equilibrium. Jurca and Faltings [8, 9] have analysis of non-truthful pure strategies. However, Jurca and Faltings [8, 9] assume the mechanism knows the prior and leave the analysis of mixed strategies as an open question. Kong et al. [11] have analysis of all equilibria including truthful and non-truthful, pure and mixed strategies while the mechanism still needs to know the prior in [11]'s setting. Our work follows the framework of BTS

and addresses the two weaknesses of BTS simultaneously when the mechanism does not know the common prior.

Now we introduce several works that have different settings than our work. The mechanisms in [4, 10] are under a different setting where agents have several a priori similar tasks. The mechanism in [4] rewards agents based on agreement which is similar to the Peer Prediction setting. That mechanism also uses the presence of multiple tasks to elicit agent strategies with high effort, and thus Dasgupta and Ghosh [4] address the equilibrium multiplicity issue for binary signals in their setting. Our setting is different since the agents only have one task (and thus we do not have to assume relations between tasks) and our results for equilibrium multiplicity issue are robust to non-binary signals. Kamble et al. [10] consider both homogeneous and heterogeneous population setting. However, this work requires large group of a priori similar tasks and their mechanisms contain equilibria that are paid higher than truth-telling. Zhang and Chen [22] propose a mechanism *Knowledge Free Peer Prediction (KFPP)* that does not know the common prior and has truth-telling as an equilibrium for small group of agents and non-binary signals. However, *Knowledge Free Peer Prediction (KFPP)* is a sequential game so agents cannot naturally report signals and predictions at the same time which is different than our setting (KFPP can be implemented non-sequentially but this requires a very complicated and unrealistic prediction reports). Finally, Cai et al. [3] have a different setting than our work while it also uses the Peer Prediction insight. The mechanism in [3] collects a set of data (x_i, y_i) to approximate a function f where y_i is a noisy version of $f(x_i)$. That mechanism rewards worker i who provides data (x_i, y_i) by comparing y_i and $\hat{f}_{-i}(x_i)$ where $\hat{f}_{-i}$ is an approximate function based on other workers' data.

Independent Work As a building block to our final mechanism, we propose a mechanism that we call the *Truthful Mechanism* which is the same as the "Divergence-Based Bayesian Truth Serum" independently proposed in [16]. In this unknown common prior mechanism, agents can naturally report their signal and prediction at the same time and truth-telling is a strict Bayesian Nash equilibrium for a small group of agents and non-binary signals without additional assumptions. However, Radanovic and Faltings [16] do not analyse non-truthful equilibria. In our paper, we analyse the structure of the equilibrium (including non-truthful equilibria) in the *Truthful Mechanism* and propose a modified mechanism that we call the *Disagreement Mechanism* which retains the same set of equilibria as the *Truthful Mechanism* yet addresses the equilibrium multiplicity issue.

2 Preliminaries, Background, and Notation

We will defer the proofs for most claims to Section 5.4.

2.1 Prior Definitions and Assumptions

We consider a setting with n agents and a set of signals Σ, and define a *setting* as a tuple (n, Σ). Each agent i has a private signal $\sigma_i \in \Sigma$ chosen from a joint distribution Q over Σ^n called the prior. Given a prior Q, for $\sigma \in \Sigma$, let $q_i(\sigma) = \Pr_Q[\sigma_i = \sigma]$ be the *a priori* probability that agent i receives signal σ. Let $q_{j,i}(\sigma'|\sigma) = \Pr_Q[\sigma_j = \sigma'|\sigma_i = \sigma]$ be the probability that agent j receives signal σ given that agent i received signal σ'.

We say that a prior Q over Σ is *symmetric* if for all σ, $\sigma' \in \Sigma$ and for all pairs of agents $i \neq j$ and $i' \neq j'$ we have $q_i(\sigma) = q_{i'}(\sigma)$ and $q_{i,j}(\sigma|\sigma') = q_{i',j'}(\sigma|\sigma')$. That is, the first two moments of the prior do not depend on the agent identities.

Assumption 2.1 (Symmetric Prior)**.** *We assume throughout that the agents' signals* $\boldsymbol{\sigma}$ *are drawn from some joint* **symmetric prior** Q*.*

Because we will assume that the prior is symmetric, we denote $q_i(\sigma)$ by $q(\sigma)$ and $q_{i,j}(\sigma|\sigma')$ (where $i \neq j$) by $q(\sigma|\sigma')$. We also define $\mathbf{q}_\sigma = q(\cdot|\sigma)$.

Assumption 2.2 (Non-zero Prior)**.** *We assume that for any* $\sigma, \sigma' \in \Sigma$*,* $q(\sigma) > 0, q(\sigma|\sigma') > 0$*.*

Assumption 2.3 (Informative Prior)**.** *We assume if agents have different private signals, they will have different expectations for the fraction of at least one signal. That is for any* $\sigma \neq \sigma'$*, there exists* σ'' *such that* $q(\sigma''|\sigma) \neq q(\sigma''|\sigma')$*.*

The following assumption conceptually states that one state is not just a more likely version of another state, and can be thought of as a weaker version of assuming $q(\sigma|\cdot)$ are linearly independent.

Assumption 2.4 (Fine-grained Prior)**.** *We assume that for any* $\sigma \neq \sigma' \in \Sigma$*, there exists* σ'', σ''' *such that*

$$\frac{q(\sigma|\sigma'')}{q(\sigma'|\sigma'')} \neq \frac{q(\sigma|\sigma''')}{q(\sigma'|\sigma''')}$$

If this assumption does not hold, then in some since σ and σ' are the same signal. We can create a new prior by replacing σ and σ' with a new signal $\sigma_0 := \sigma$ or σ', and not lose any information, in the sense that we can still recover the original prior. To see this, we first define $p = \frac{q(\sigma)}{q(\sigma')}$, and note that for all σ'', $p = \frac{q(\sigma|\sigma'')}{q(\sigma'|\sigma'')}$. Whenever σ_0 is drawn in the new prior, we simply replace it by σ with probability p and σ' with probability $1 - p$. This produces the same prior for agents that have no information or other their signal's information.

We illustrate this in the below example:

Example 2.5. $Q = \begin{pmatrix} q(s_1|s_1) & q(s_1|s_2) & q(s_1|s_3) \\ q(s_2|s_1) & q(s_2|s_2) & q(s_2|s_3) \\ q(s_3|s_1) & q(s_3|s_2) & q(s_3|s_3) \end{pmatrix} = \begin{pmatrix} 0.1 & 0.2 & 0.3 \\ 0.2 & 0.4 & 0.6 \\ 0.7 & 0.4 & 0.1 \end{pmatrix}$ *is not a fine-grained prior since*

$$\frac{q(s_1|s_1)}{q(s_2|s_1)} = \frac{q(s_1|s_2)}{q(s_2|s_2)} = \frac{q(s_1|s_3)}{q(s_2|s_3)}$$

Note that in this example, even we combine s_1 and s_2 to be a single signal s_0 which is defined as $s_0 := s_1$ or s_2, we do not lose any information: if an agent knows that the fraction of agents who report s_0 is x, we know his belief for the expectation of the fraction of s_1 must be $\frac{x}{3}$ no matter what private signal he receives.

We only require the fine-grained prior assumption to show that truth-telling is *strictly* "better" than any other symmetric equilibrium (excluding permutation equilibrium). In the above example where the prior is not fine-grained, if agents always report s_1 when they receive s_1 or s_2, this does not lose information (is not "worse") comparing with the case agents always tell the truth. So we cannot say truth-telling is strictly "better" than any other equilibrium when the prior is not fine-grained. However, this assumption is not necessary to show that truth-telling is a strict Bayesian equilibrium of our mechanism, nor to show that the agent welfare of truth-telling is at least as high as other symmetric equilibrium.

Assumption 2.6 (Ensemble Prior)**.** *Although we talk of a single prior, in fact we have an ensemble* $Q = \{Q_n\}_{n \in N, n \geq 3}$ *of priors; one for each possible number of agents greater than 3. We assume that all* Q_n *are over the same signal set* Σ *have have identical* $q(\sigma)$ *and* $q(\sigma'|\sigma)$*.*

When the number of agents n changes, the joint prior actually changes as well, but the first two moments of the prior are fixed. This allows us to make meaningful statements about n going to infinity.

We sometimes will denote the class of priors that satisfy all five of these assumptions as SNIFE priors.

2.2 Game Setting and Equilibrium Concepts

Given a setting (n, Σ) with prior Q, we consider a game in which each agent i is asked to report his private signal $\sigma_i \in \Sigma$ and his prediction $\mathbf{p}_i \in \Delta_\Sigma$, a distribution over Σ, where $\mathbf{p}_i = \mathbf{q}_{\sigma_i}$. For any $\sigma \in \Sigma$, $\mathbf{p_i}(\sigma)$ is agent i's (reported) expectation for the fraction of other agents who has received σ given he has received σ_i. However, agents may not tell the truth. We denote $\Sigma \times \Delta_\Sigma$ by $\mathcal{R}$. We define a report profile of agent i as $r_i = (\hat{\sigma}_i, \hat{\mathbf{p}}_i) \in \mathcal{R}$ where $\hat{\sigma}_i$ is agent i's reported signal and $\hat{\mathbf{p}}_i$ is agent i's reported prediction.

We would like to encourage truth-telling, namely that agent i reports $\hat{\sigma}_i = \sigma_i, \hat{\mathbf{p}}_i = \mathbf{q}_{\sigma_i}$. To this end, agent i will receive some payment $\nu_i(\hat{\sigma}_i, \hat{\mathbf{p}}_i, \hat{\sigma}_{-i}, \hat{\mathbf{p}}_{-i})$ from our mechanism.

Now we consider the strategy an agent plays in the game.

Definition 2.7 (Strategy)**.** Given a mechanism $\mathcal{M}$, we define the strategy of $\mathcal{M}$ for setting (n, Σ) as a mapping s from (σ, Q) (the signal and common prior received) to a probability distribution over $\mathcal{R}$ (the reported signal, prediction pair).

That is, for each possible signal σ and prior Q an receives, he will choose a signal, prediction pair to report from some distribution $s(\sigma, Q)$. We define a strategy profile $\mathbf{s}$ as a profile of all agents' strategies $\{s_1, s_2, ...s_n\}$ and we say agents play $\mathbf{s}$ if for any i, agent i plays strategy s_i. We say a strategy profile is **symmetric** if each agent plays the same strategy.

We define the **agent welfare** of a strategy profile $\mathbf{s}$ and a mechanism $\mathcal{M}$ for setting (n, Σ) with prior Q to be the expectation of the sum of payments to each agent and we write it as $AW_{\mathcal{M}}(n, \Sigma, Q, \mathbf{s})$.

A *Bayesian Nash equilibrium* consists of a strategy profile $s = (s_1, \ldots, s_n)$ such that no player wihes to change her strategy, given the strategies of the other players and the information contained in the prior and her signal. Formally,

Definition 2.8 (Bayesian Nash equilibrium)**.** Given a family of priors $\mathcal{Q}$, a strategy profile $s = (s_1, \ldots, s_n)$ is a Bayesian Nash equilibrium if and only if for any prior $Q \in \mathcal{Q}$, for any i, and for any s_i'

$$\begin{aligned}&\mathbb{E}_{(\hat{\sigma}_i', \hat{\mathbf{p}}_i') \leftarrow s_i'(\sigma_i, Q), (\hat{\sigma}_{-i}, \hat{\mathbf{p}}_{-i}) \leftarrow s_{-i}(\sigma_{-i}, Q)}[\nu_i(\hat{\sigma}_i', \hat{\mathbf{p}}_i', \hat{\sigma}_{-i}, \hat{\mathbf{p}}_{-i})] \\ \leq &\mathbb{E}_{(\hat{\sigma}_i, \hat{\mathbf{p}}_i) \leftarrow s_i(\sigma_i, Q), (\hat{\sigma}_{-i}, \hat{\mathbf{p}}_{-i}) \leftarrow s_{-i}(\sigma_{-i}, Q)}[\nu_i(\hat{\sigma}_i, \hat{\mathbf{p}}_i, \hat{\sigma}_{-i}, \hat{\mathbf{p}}_{-i})]\end{aligned}$$

In the case where, for some i, the equality holds if and only if $s_i' = s_i$, we say this strategy profile is a *strict Bayesian Nash equilibrium* for prior family $\mathcal{Q}$.

Remark 2.9 (Equilibrium for a Given Prior)**.** *Note that we assume agents have a common prior Q, so often for convenience, we will implicitly assume Q is fixed, at which point a strategy is a mapping from Σ to a probability distribution over $\mathcal{R}$. We will call such a strategy profile $\mathbf{s}$ an equilibrium for prior Q if it satisfies the condition of Bayesian Nash equilibrium when Q is fixed.*

Assuming a fixed prior Q, for any strategy profile $s = (s_1, s_2, ..., s_n)$, we will represent the marginal distribution of an agent i's strategy for her signal report as a matrix θ_i where $\theta_i(\hat{\sigma}, \sigma)$

is the probability that agent will report signal $\hat{\sigma}$ when his private signal is σ. Note that θ_i is a **transition matrix**, that is the sum of every column is 1. We call θ_i the signal strategy of agent i. We also call $(\theta_1, \theta_2, ..., \theta_n)$ the signal strategy of s. We define the *average signal strategy* of s as $\bar{\theta}_n = \frac{\sum_i \theta_i}{n}$. The following claim relates this average signal strategy to the distribution of all reported signals:

Claim 2.10. *Assume that the distribution over all agents' private signals is $\omega \in \Delta_\Sigma$, the distribution over all agents' reported signals will be $\bar{\theta}_n \omega$.*

Note that the mechanism actually collects agents' reported signals, so in order to estimate the distribution over their private signals, we hope $\bar{\theta}_n$ is (close to) the identity matrix I.

2.3 Special Strategy Profiles

In this section, we will introduce three special types of strategy profiles that we call *truth-telling*, *best prediction strategy profiles*, and *permutation strategy profiles*.

Definition 2.11 (Truth-telling)**.** We define a strategy profile as truth-telling if for all i, and for all Q, $s(\sigma_i, Q) = (\sigma_i, \mathbf{q}_{\sigma_i})$ with probability 1. We write the truth-telling strategy profile as $\mathbf{T}$.

For every agent i, let $\hat{\sigma}$ be a randomly chosen agent's reported signal, when other agents tell the truth, the distribution of $\hat{\sigma}$ is $\mathbf{q}_{\sigma_i}$. However, if agents play strategy $\mathbf{s}$, for agent i, the distribution of $\hat{\sigma}$ depends on not only his prior Q but also the strategy $\mathbf{s}$. We define the distribution of $\hat{\sigma}$ for agent i as $\mathbf{q}^{\mathbf{s}}_{\sigma_i}$.

Claim 2.12.

$$\mathbf{q}^{\mathbf{s}}_{\sigma_i} = \theta_{-i} \mathbf{q}_{\sigma_i}$$

where $(\theta_1, \theta_2, ..., \theta_n)$ is $\mathbf{s}$'s signal strategy and $\theta_{-i} = \frac{\sum_{j \neq i} \theta_j}{n-1}$.

When agents play strategy $\mathbf{s}$, to best predict other agents' reported signal, agent i should be report $\mathbf{q}^{\mathbf{s}}_{\sigma_i}$ rather than $\mathbf{q}_{\sigma_i}$. This motivates our definition for *best prediction strategy profile* which is a strategy profile where every agent i gives his "best prediction" $\mathbf{q}^{\mathbf{s}}_{\sigma_i}$.

Definition 2.13 (Best Prediction Strategy Profile)**.** We say a strategy profile $\mathbf{s}$ is a best prediction strategy profile if for every agent i, he reports $\mathbf{q}^{\mathbf{s}}_{\sigma_i}$. We call a best strategy prediction strategy profile $\mathbf{s}$ a *symmetric best strategy prediction strategy profile* if $\theta_i = \theta$ for every i.

Now we begin to introduce the definition of a permutation strategy profile. Intuitively, if agents "collude" to relabel the signals and then tell the truth with relabeled signals, they actually play what we will call permutation strategy profile.

Given a permutation $\pi : \Sigma \mapsto \Sigma$ (which is actually a relabeling of signals), by abusing notation a little bit, we define $\pi : \mathcal{Q} \mapsto \mathcal{Q}$ as a mapping from a prior Q to a *permuted prior* $\pi(Q)$ where for any $\sigma_1, \sigma_2, ..., \sigma_n \in \Sigma$,

$$Pr_{\pi(Q)}(\sigma_1, \sigma_2, ..., \sigma_n) = Pr_Q(\pi^{-1}(\sigma_1), \pi^{-1}(\sigma_2), ..., \pi^{-1}(\sigma_n))$$

where σ_i is the private signal of agent i. Notice that it follows that:

$$Pr_{\pi(Q)}(\pi(\sigma_1), \pi(\sigma_2), ..., \pi(\sigma_n)) = Pr_Q(\sigma_1, \sigma_2, ..., \sigma_n).$$

Intuitively, $\pi(Q)$ is the same with Q when the signals are relabeled according to π.

For any strategy s, we define $\pi(s)$ as the strategy such that $\pi(s)(\sigma, Q) = s(\pi(\sigma), \pi(Q))$.

Definition 2.14 (Permuted Strategy Profile)**.** For any strategy profile $\mathbf{s}$, we define $\pi(\mathbf{s})$ as a strategy profile with $\pi(\mathbf{s}) = (\pi(s_1), \pi(s_2), ..., \pi(s_n))$.

Note that $\pi^{-1}\pi Q = Q$ which implies $\pi^{-1}\pi(\mathbf{s}) = \mathbf{s}$.

Definition 2.15 (Permutation Strategy Profile)**.** We define a strategy profile $\mathbf{s}$ as a permutation strategy profile if there exists a permutation $\pi : \Sigma \to \Sigma$ such that $\mathbf{s} = \pi(\mathbf{T})$.

Note that if agents play $\pi(\mathbf{T})$, then the signal strategy of each agent is π, and so the distribution of report profiles is $\bar{\theta}_n\omega = \pi\omega$.

There exists a natural bijection between permutation strategy profiles and $|\Sigma| \times |\Sigma|$ permutation matrices. If the permutation strategy profile is constructed by permutation π, the only non-zero entries of the corresponding permutation matrix θ_π are $\theta_\pi(\pi(\sigma), \sigma) = 1$ for all $\sigma \in \Sigma$. For a transition matrix θ, if θ is not a permutation matrix, we would like to give a definition for when a transition matrix θ is what we call τ-*close* to a permutation given any $\tau > 0$. This definition is motivated by the below claim and will be described after it.

Claim 2.16. *For any transition matrix $\theta_{m\times m}$ where the sum of every column is 1, θ is a permutation matrix iff for any row of θ, there at most one non-zero entry.*

Now we give a definition for τ-*close*.

Definition 2.17 (τ-close)**.** We say a signal strategy θ is τ-*close* to a permutation if for any row of θ, there is at most one entry that is greater than τ.

2.4 Mechanism Design Goals

We want to design "good" mechanisms that motivate the agents to reveal their private information truthfully. Now we will give several definitions for "good" mechanisms:

We say that a mechanism is ***truthful*** if truth-telling is a strict Bayesian Nash equilibrium whenever Q is a symmetric and informative distribution.

We say that a mechanism has truth-telling as a ***focal*** equilibrium if, for any SNIFE prior Q, the mechanism is truthful and the agent welfare is strictly higher in the truth-telling equilibrium than in any other any other Bayesian Nash equilibrium. Recall that the agent welfare is the expected sum of payments for all agents.

However, it turns out that making truth-telling *focal* is too much to ask. We will show that with an unknown common prior, for any mechanism, and any permutation strategy profile, there exists a prior such that that permutation strategy profile is an equilibrium and has agent welfare at least as much as truth-telling. One natural question is whether there is any mechanism such that all permutation strategy profiles have agent welfare equal to each other, but strictly higher than any other Bayesian Nash equilibrium? This question motivates a weaker version of *focal*: *quasi-focal*.

Formally, we say that a mechanism has truth-telling as a ***quasi-focal*** equilibrium if, for any SNIFE prior Q, the mechanism is truthful and agent-welfare is strictly higher in the permutation equilibrium than in any other Bayesian Nash equilibrium where the agents do not play a permutation strategy profile.

It will turn out that the mechanism we purpose cannot make truth-telling *quasi-focal*. However, we can show that it satisfies the following three slight relaxations of the definition:

***Symmetric-quasi-focal*:** We say truth-telling is symmetric-quasi-focal in a mechanism if

1. all permutation equilibrium have equal agent welfare; and

2. any *symmetric* equilibrium that is not a permutation equilibrium has agent-welfare strictly less than truth-telling.

***(τ_1, γ_1)-robust-symmetric-quasi-focal*:** We say truth-telling is *(τ_1, γ_1)-robust-symmetric-quasi-focal* in a mechanism if any *symmetric* equilibrium that pays within γ_1 of truth-telling must be $\tau_1(\gamma_1)$ close to a permutation strategy profile.

***(τ_2, γ_2)-robust-approximate-quasi-focal*:** We say truth-telling is *(τ_2, γ_2)-robust-approximate-quasi-focal* in a mechanism if

1. all permutation equilibrium pay the same;
2. no equilibrium has agent welfare greater than $\gamma_2(n)$ more than that of truth-telling where n is the number of agents; and
3. any profile that pays within $\gamma_2(n)$ of truth-telling must be $\tau_2(n)$ close to a permutation strategy profile.

2.5 F-divergence

Now we introduce f-divergence, a key tool we will use in our mechanism design. f-divergence([2]) is used to measuring the "difference" between distributions. One important property of f-divergence is information monotonicity: For any two distributions, if we post-process each distribution in the same way, the two distributions will become "closer" because of the information loses.

***F*-divergence** [1] $D_f : \Delta_\Sigma \times \Delta_\Sigma \to \mathbb{R}$ is a non-symmetric measure of difference between a distribution $\mathbf{p} \in \Delta_\Sigma$ and a distribution $\mathbf{q} \in \Delta_\Sigma$ and is defined to be

$$D_f(\mathbf{p}, \mathbf{q}) = \sum_{\sigma \in \Sigma} \mathbf{p}(\sigma) f\left(\frac{\mathbf{p}(\sigma)}{\mathbf{q}(\sigma)}\right)$$

where $f(\cdot)$ is a convex function. Now we introduce the properties of f-divergence:

1. **Non-negative**: For any $\mathbf{p}, \mathbf{q}$, $D_f(\mathbf{p}, \mathbf{q}) \geq 0$ and $D_f(\mathbf{p}, \mathbf{q}) = 0$ if and only if $\mathbf{p} = \mathbf{q}$.
2. **Convexity**: Both $D_f(\cdot, \mathbf{q})$ and $D_f(\mathbf{p}, \cdot)$ are convex functions for any $\mathbf{p}, \mathbf{q}$.
3. **Information Monotonicity**: For any $\mathbf{p}, \mathbf{q}$, and transition matrix $\theta \in \mathbb{R}^{|\Sigma| \times |\Sigma|}$ where $\theta(\sigma, \sigma')$ is the probability that we map σ' to σ, we have $D_f(\mathbf{p}, \mathbf{q}) \geq D_f(\theta\mathbf{p}, \theta\mathbf{q})$. When θ is a permutation, $D_f(\mathbf{p}, \mathbf{q}) = D_f(\theta\mathbf{p}, \theta\mathbf{q})$.

Now we introduce the proof in ([1]) for information monotonicity and give the conditions for the inequality of information monotonicity to be strict.

Lemma 2.18 (Information Monotonicity ([1]))**.** *For any strictly convex function f, f-divergence $D_f(\mathbf{p}, \mathbf{q})$ satisfies information monotonicity so that for any transition matrix $\theta \in \mathbb{R}^{\Sigma \times \Sigma}$, $D_f(\mathbf{p}, \mathbf{q}) \geq D_f(\theta\mathbf{p}, \theta\mathbf{q})$.*

Moreover, the inequality is strict if and only if there exists $\sigma, \sigma', \sigma''$ such that $\theta(\sigma, \sigma')\mathbf{p}(\sigma') > 0$, $\theta(\sigma, \sigma'')\mathbf{p}(\sigma'') > 0$ and $\frac{\mathbf{p}(\sigma'')}{\mathbf{p}(\sigma')} \neq \frac{\mathbf{q}(\sigma'')}{\mathbf{q}(\sigma')}$.

We give an example where the strictness condition is not satisfied in appendix.

Proof. The proof follows from algebraic manipulation and one application of convexity.

$$D_f(\theta\mathbf{p}, \theta\mathbf{q}) = \sum_\sigma (\theta\mathbf{p})(\sigma) f\left(\frac{(\theta\mathbf{q})(\sigma)}{(\theta\mathbf{p})(\sigma)}\right) \tag{1}$$

$$= \sum_\sigma \theta(\sigma, \cdot)\mathbf{p} f\left(\frac{\theta(\sigma, \cdot)\mathbf{q}}{\theta(\sigma, \cdot)\mathbf{p}}\right) \tag{2}$$

$$= \sum_\sigma \theta(\sigma, \cdot)\mathbf{p} f\left(\frac{1}{\theta(\sigma, \cdot)\mathbf{p}} \sum_{\sigma'} \theta(\sigma, \sigma')\mathbf{p}(\sigma')\frac{\mathbf{q}(\sigma')}{\mathbf{p}(\sigma')}\right) \tag{3}$$

$$\leq \sum_\sigma \theta(\sigma, \cdot)\mathbf{p} \frac{1}{\theta(\sigma, \cdot)\mathbf{p}} \sum_{\sigma'} \theta(\sigma, \sigma')\mathbf{p}(\sigma') f\left(\frac{\mathbf{q}(\sigma')}{\mathbf{p}(\sigma')}\right) \tag{4}$$

$$= \sum_\sigma \mathbf{p}(\sigma) f\left(\frac{\mathbf{q}(\sigma)}{\mathbf{p}(\sigma)}\right) = D_f(\mathbf{p}, \mathbf{q}) \tag{5}$$

The second equality holds since $(\theta\mathbf{p})(\sigma)$ is dot product of the σ^{th} row of θ and $\mathbf{p}$.
The third equality holds since $\sum_{\sigma'} \theta(\sigma, \sigma')\mathbf{p}(\sigma')\frac{\mathbf{q}(\sigma')}{\mathbf{p}(\sigma')} = \theta(\sigma, \cdot)\mathbf{q}$.
The fourth inequality follows from the convexity of $f(\cdot)$.
The last equality holds since $\sum_\sigma \theta(\sigma, \sigma') = 1$.

We now examine under what conditions the inequality in Equation 4 is strict. Note that for any strictly convex function g, if $\forall u, \lambda_u > 0$, $g(\sum_u \lambda_u x_u) = \sum_u \lambda_u g(x_u)$ if and only if there exists x such that $\forall u, x_u = x$. By this property, the inequality is strict if and only if there exists $\sigma, \sigma', \sigma''$ such that $\frac{\mathbf{p}(\sigma'')}{\mathbf{p}(\sigma')} \neq \frac{\mathbf{q}(\sigma'')}{\mathbf{q}(\sigma')}$ and $\theta(\sigma, \sigma')\mathbf{p}(\sigma') > 0$, $\theta(\sigma, \sigma'')\mathbf{p}(\sigma'') > 0$.

□

Corollary 2.19. *Given SNIFE prior Q, for any θ that is not a permutation, there exists two private signals $\sigma_1 \neq \sigma_2$ such that $D_f(\theta\mathbf{q}_{\sigma_1}, \theta\mathbf{q}_{\sigma_2}) < D_f(\mathbf{q}_{\sigma_1}, \mathbf{q}_{\sigma_2})$*

Proof. First notice that when θ is not a permutation, based on Claim 2.16, there exists a row of θ such that the row has at least two positive entries, in other words, there exists $\sigma, \sigma', \sigma''$ such that $\theta(\sigma, \sigma'), \theta(\sigma, \sigma'') > 0$. Based on the non-zero and fine-grained assumptions of Q, there exists $\sigma_1 \neq \sigma_2$ such that
$\theta(\sigma, \sigma')\mathbf{p}(\sigma'), \theta(\sigma, \sigma'')\mathbf{p}(\sigma'') > 0$ and $\frac{\mathbf{p}(\sigma')}{\mathbf{p}(\sigma'')} \neq \frac{\mathbf{q}(\sigma')}{\mathbf{q}(\sigma'')}$ where $\mathbf{p} = \mathbf{q}_{\sigma_1}, \mathbf{q} = \mathbf{q}_{\sigma_2}$. When
$\theta(\sigma, \sigma')\mathbf{p}(\sigma'), \theta(\sigma, \sigma'')\mathbf{p}(\sigma'') > 0$, we have $\theta(\sigma, \cdot)\mathbf{p} > 0$. By Lemma 2.18, we have $D_f(\theta\mathbf{q}_{\sigma_1}, \theta\mathbf{q}_{\sigma_2}) < D_f(\mathbf{q}_{\sigma_1}, \mathbf{q}_{\sigma_2})$ □

Now we will introduce a special f-divergence called Hellinger-divergence and then we will give several properties of Hellinger-divergence ([1]).

Hellinger-divergence If we pick the convex function $f(\cdot)$ to be $(\sqrt{x} - 1)^2$, we will obtain Hellinger-divergerce

$$D^*(\mathbf{p}, \mathbf{q}) = \sum_\sigma (\sqrt{\mathbf{p}(\sigma)} - \sqrt{\mathbf{q}(\sigma)})^2$$

Since Hellinger-divergence is an f-divergence, it also has convexity and information monotonicity. Besides these two properties, Hellinger-divergence has another two properties which will be used in the future.

1. **Bounded divergence**: $0 \leq D^*(\mathbf{p}, \mathbf{q}) \leq 1$

2. **Square root triangle inequality**: $\sqrt{D^*(\cdot,\cdot)}$ is a metric.

Note that Hellinger-divergence is bounded which is different than KL-divergence. Since we use Hellinger-divergence in our disagreement mechanism, we can always guarantee bounded payment.

2.6 Proper Scoring Rules

Now we introduce strictly proper scoring rules, another key tool we will use in our mechanism design. Starting with [12], proper scoring rules have become a common ingredient in mechanisms for elicit unverifiable information elicitation (e.g. [14, 18]).

A scoring rule $PS : \Sigma \times \Delta_\Sigma \to \mathbb{R}$ takes in a signal $\sigma \in \Sigma$ and a distribution over signals $\delta_\Sigma \in \Delta_\Sigma$ and outputs a real number. A scoring rule is *proper* if, whenever the first input is drawn from a distribution δ_Σ, then the expectation of PS is maximized by δ_Σ. A scoring rule is called *strictly proper* if this maximum is unique. We will assume throughout that the scoring rules we use are strictly proper. By slightly abusing notation, we can extend a scoring rule to be $PS : \Delta_\Sigma \times \Delta_\Sigma \to \mathbb{R}$ by simply taking $PS(\delta_\Sigma, \delta'_\Sigma) = \mathbb{E}_{\sigma \leftarrow \delta_\Sigma}(\sigma, \delta'_\Sigma)$. We note that this means that any proper scoring rule is linear in the first term.

Example 2.20 (Example of Proper Scoring Rule)**.** *Fix an outcome space Σ for a signal σ. Let $\mathbf{q} \in \Delta_\Sigma$ be a reported distribution. The Logarithmic Scoring Rule maps a signal and reported distribution to a payoff as follows:*

$$L(\sigma, \mathbf{q}) = \log(\mathbf{q}(\sigma)).$$

Let the signal σ be drawn from some random process with distribution $\mathbf{p} \in \Delta_\Sigma$.
Then the expected payoff of the Logarithmic Scoring Rule

$$\mathbb{E}_{\sigma \leftarrow \mathbf{p}}[L(\sigma, \mathbf{q})] = \sum_\sigma \mathbf{q}(\sigma) \log \mathbf{q}(\sigma) = L(\mathbf{p}, \mathbf{q})$$

According to [7], this value will be maximized if and only if $\mathbf{q} = \mathbf{p}$.

3 Impossibility of Truth-telling Being Focal

In this section, we show the impossibility of truth-telling being focal. In fact, we show something stronger, that no strategy profile can always be the equilibrium and have agent welfare that is strictly greater than any other equilibrium (including truth-telling). Moreover, our impossibility result applies to a very general setting of mechanisms. The case that we consider, where each agent reports a signal and prediction pair, is a special case of this general setting.

We first define the class of mechanisms to which our impossibility result will apply.

Definition 3.1 (Mechanism)**.** We define a mechanism $\mathcal{M}$ for a setting (n, Σ) as a tuple $\mathcal{M} := (\mathcal{R}, M)$ where $\mathcal{R}$ is a set of all possible report profiles the mechanism allows, and $M : \mathcal{R}^n \mapsto \mathbb{R}^n$ is a mapping from all agents' report profiles to each agent's reward.

The intuitive explanation for this impossibility result is that the agents can collude to relabel the signals and the mechanism has no way to defend against this relabeling without knowing some information about agents' common prior.

The proposition stated below implies that, in particular, no strategy profile can always be an equilibrium that has the agent welfare that is strictly greater than any other equilibrium.

Proposition 3.2. *Let $\mathcal{M}$ be a mechanism that does not know the common prior, for any strategy profile s, and any permutation π:*
(1) s is a strict Bayesian Nash equilibrium of $\mathcal{M}$ for any symmetric, informative prior iff $\pi(s)$ is a strict Bayesian Nash equilibrium of $\mathcal{M}$ for any symmetric, informative prior.
(2) There exists a prior Q such that $AW_{\mathcal{M}}(n, \Sigma, Q, \mathbf{s}) \leq AW_{\mathcal{M}}(n, \Sigma, Q, \pi(\mathbf{s}))$.

The key idea to prove this theorem is what we refer to as **Indistinguishable Scenarios**:

Definition 3.3 (Scenario)**.** We define a scenario for the setting (n, Σ) as a tuple $(Q, \mathbf{s})$ where Q is a prior, and $\mathbf{s}$ is a strategy profile.

Given mechanism $\mathcal{M}$, for any scenario $A = (Q_A, \mathbf{s}_A)$, we write $AW_{\mathcal{M}}(n, \Sigma, A)$ as the agent welfare when agents play $\mathbf{s}_A$ and have common prior Q_A.

For two scenarios $A = (Q_A, \mathbf{s}_A)$, $B = (Q_B, \mathbf{s}_B)$ for setting (n, Σ), let $\sigma_A := (\sigma_{1_A}, \sigma_{2_A}, ..., \sigma_{n_A})$ be agents' private signals drawn from Q_A, $\sigma_B := (\sigma_{1_B}, \sigma_{2_B}, ..., \sigma_{n_B})$ be agents' private signals drawn from Q_B.

Definition 3.4 (Indistinguishable Scenarios)**.** We say two scenarios A, B are indistinguishable $A \approx B$ if there is a coupling of the random variables σ_A and σ_B such that $\forall i$, $s_A(\sigma_{i_A}, Q_A) = s_B(\sigma_{i_B}, Q_B)$ and agent i_A has the same belief about the world as agent i_B, in other words, $Pr(\hat{\sigma}|\sigma_{i_A}, Q_A, s_A) = Pr(\hat{\sigma}|\sigma_{i_B}, Q_B, s_B)\ \forall \hat{\sigma} \in \Sigma$.

Now we will prove two properties of indistinguishable scenarios which are the main tools in the proof for our impossibility result.

Observation 3.5. *If $(Q_A, \mathbf{s}_A) \approx (Q_B, \mathbf{s}_B)$, then (i) for any mechanism $\mathcal{M}$, $\mathbf{s}_A$ is a (strict) equilibrium for prior Q_A iff $\mathbf{s}_B$ is a (strict) equilibrium for prior Q_B. (ii) $AW_{\mathcal{M}}(n, \Sigma, A) = AW_{\mathcal{M}}(n, \Sigma, B)$*

At a high level, (1) is true since any reported profile distribution that agent i_A can deviate to, agent i_B can deviate to the same reported profile distribution as well and obtain the same expected payment as agent i_A.

Formally, we will prove the $\Rightarrow$ direction in (1) by contradiction. The proof of the other direction will be similar. Consider the coupling for σ_A, σ_B mentioned in the definition of indistinguishable scenarios. For the sake of contradiction, assume there exists i and σ_{i_B} such that $\hat{\sigma}' \neq s_{i_B}(\sigma_{i_B}, Q_B)$ is a best response for agent i_B. Since agent i_A has the same belief about the world as agent i_B and $s_{i_A}(\sigma_{i_A}, Q_A) = s_{i_B}(\sigma_{i_B}, Q_B)$, $\hat{\sigma}' \neq s_{i_A}(\sigma_{i_A}, Q_A)$ is a best response to agent i_A as well, which is a contradiction to the fact that $\mathbf{s}_A$ is a strictly equilibrium for prior Q_A.

To gain intuition about (2), consider the coupling again. For any i, agent i_A reports the same thing and has the same belief for the world as agent i_B, which implies the expected payoff of agent i_A is the same as agent i_B. (2) follows.

Now we are ready to prove our impossibility result:

of Proposition 3.2. We prove part (1) and part (2) separately.

Proof of Part (1) Let $A := (Q, \mathbf{s}), B := (\pi^{-1}(Q), \pi(\mathbf{s}))$. We will show that for any strategy profile $\mathbf{s}$ and any prior Q, $A \approx B$. Based on our above observations, part (1) immediately follows from that fact.

To prove $(Q, \mathbf{s}) \approx (\pi^{-1}Q, \pi(\mathbf{s}))$, we can couple $(\sigma_1, \sigma_2, ..., \sigma_n)$ with $(\pi^{-1}(\sigma_1), \pi^{-1}(\sigma_2), .., \pi^{-1}(\sigma_n))$ where $(\sigma_1, \sigma_2, ..., \sigma_n)$ is drawn from Q. It is a legal coupling since

$$Pr_{\pi^{-1}(Q)}(\pi^{-1}(\sigma_1), \pi^{-1}(\sigma_2), .., \pi^{-1}(\sigma_n)) = Pr_Q(\sigma_1, \sigma_2, ..., \sigma_n)$$

according to the definition of $\pi^{-1}(Q)$.

Now we show this coupling satisfies the condition in Definition 3.4. First note that $\pi(s_i)(\pi^{-1}(\sigma_i), \pi^{-1}(Q)) = s_i(\sigma_i, Q)$. Now we begin to calculate $Pr(\hat{\sigma}|\sigma_{i_B}, Q_B, s_B)$

$$\begin{aligned}
Pr(\hat{\sigma}|\sigma_{i_B}, Q_B, s_B) =& Pr(\hat{\sigma}|\pi^{-1}\sigma_i, \pi^{-1}(Q), \pi(s)) && (6)\\
=& \sum_j \Pr(j) \sum_{\sigma'} Pr_{\pi^{-1}(Q)}(\sigma'|\pi^{-1}\sigma_i) Pr(\pi(s_j)(\sigma', \pi^{-1}(Q)) = \hat{\sigma}) && (7)\\
=& \sum_j \Pr(j) \sum_{\sigma'} Pr_{\pi^{-1}(Q)}(\sigma'|\pi^{-1}\sigma_i) Pr(s_j(\pi(\sigma'), \pi\pi^{-1}(Q)) = \hat{\sigma}) && (8)\\
=& \sum_j \Pr(j) \sum_{\sigma'} Pr_Q(\pi(\sigma')|\sigma_i) Pr(s_j(\pi(\sigma'), Q) = \hat{\sigma}) && (9)\\
=& \sum_j \Pr(j) \sum_{\sigma''} Pr_Q(\sigma''|\sigma_i) Pr(s_j(\sigma'', Q) = \hat{\sigma}) && (10)\\
=& Pr(\hat{\sigma}|\sigma_i, Q, s) = Pr(\hat{\sigma}|\sigma_{i_A}, Q_A, s_A) && (11)
\end{aligned}$$

From (6) to (7): To calculate a randomly chosen agent's reported signal, we should sum over all possible agents j and calculate the probability conditioning on agent j being picked. To calculate the probability that agent j has reported $\hat{\sigma}$, we should sum over all possible private signals agent j has received and calculate the probability agent j reported $\hat{\sigma}$ conditioning on he received private signal σ', which is determined by agent j's strategy.

By abusing notation a little bit, we can write $\pi(s_j)(\sigma', \pi^{-1}Q)$ as a random variable (it is actually a distribution) with $Pr(\pi(s_j)(\sigma', \pi^{-1}Q) = \hat{\sigma}) = \pi(s_j)(\sigma', \pi^{-1}Q)(\hat{\sigma})$. According to above explanation, (7) follows.

(8) follows from the definition of permuted strategy (See Section 2.3).

(9) follows from the definition of permuted prior (See Section 2.3).

By replacing $\pi(\sigma')$ by σ'', (10) follows.

We finished the proof $A \approx B$, as previously argued, result (1) follows.

Proof for Part (2) We will prove the second part by contradiction:

Fix permutation strategy profile π. First notice that there exists an positive integer O_d such that $\pi^{O_d} = I$ where I is the identity and agents play I means they tell the truth (we can pick O_d as the order of π in the permutation group).

Given any strategy profile s, for the sake of contradiction, we assume that there exists a mechanism $\mathcal{M}$ with unknown common prior such that $AM_{\mathcal{M}}(n, \Sigma, Q, \mathbf{s}) > AM_{\mathcal{M}}(n, \Sigma, Q, \pi(\mathbf{s}))$ for any prior Q. For positive integer $k \in \{0, 1, ..., O_d\}$, we construct three scenarios:

$$A_k := (\pi^k(Q), s),\ A_{k+1} := (\pi^{k+1}(Q), s),\ B_k := (\pi^k(Q),\ \pi(s))$$

and show for any k,

(I)$AM_{\mathcal{M}}(n, \Sigma, A_k) > AM_{\mathcal{M}}(n, \Sigma, B_k)$,

(II) $AM_{\mathcal{M}}(n, \Sigma, A_{k+1}) = AM_{\mathcal{M}}(n, \Sigma, B_k)$.

Combining (I), (II) and the fact $A_0 = A_{O_d}$, we have

$$AM_{\mathcal{M}}(n, \Sigma, A_0) > AM_{\mathcal{M}}(n, \Sigma, A_1) > ...AM_{\mathcal{M}}(n, \Sigma, A_{O_d}) = AM_{\mathcal{M}}(n, \Sigma, A_0)$$

which is a contradiction.

Now it is only left to show (I) and (II). Based on our assumption

$$AM_{\mathcal{M}}(n, \Sigma, Q, \mathbf{s}) > AM_{\mathcal{M}}(n, \Sigma, Q, \pi(\mathbf{s}))$$

for any prior Q, we have (I). By the same proof we have in part (1), we have $A_{k+1} \approx B_k$, which implies (II) according to our above observations.

□

Corollary 3.6. *Let $\mathcal{M}$ be a mechanism that does not know the common prior, given truth-telling strategy $\mathbf{T}$, if there exists a permutation π such that $\pi(\mathbf{T}) \neq \mathbf{T}$, $\mathbf{T}$ cannot have agent welfare that is always strictly highest among all equilibria.*

The requirement that $\pi(\mathbf{T}) \neq \mathbf{T}$ ensures, that a truthful input to the mechanism depends on the private signal of an agent. An example when this would not hold would be if there is only possible input.

Proof. According to Proposition 3.2, $\pi(\mathbf{T})$ is an equilibrium as well and there exists a prior Q such that the agent welfare of $\pi(\mathbf{T})$ is greater or equal to $\mathbf{T}$, so $\mathbf{T}$ cannot have agent welfare that is always strictly highest among all equilibria. □

4 The Disagreement Mechanism

In this section, we will introduce two mechanisms: the *Truthful Mechanism* and the *Disagreement Mechanism*. Both of these mechanisms are based on the Bayesian Truth serum (BTS) framework, which means agents are paid based on a "prediction score" and an "information score". We will introduce the "prediction score" used in the two mechanisms first, it will become a strong tool to motivate people to tell the truth.

Prediction Score via Proper Scoring Rules Agents will receive a prediction score based on how well their prediction predicts a randomly chosen agent's reported signal. Say an agent i reports prediction $\hat{\mathbf{p}}_i$ then a random agent, call him agent j, is picked for him, agent i will receive a prediction score $PS(\hat{\sigma}_j, \hat{\mathbf{p}}_i)$ when PS is a proper scoring rule (see Section 2.5) of proper scoring rules family. Note that any proper scoring rule works. $PS(\hat{\sigma}_j, \hat{\mathbf{p}}_i)$ is maximized if and only if agent i's reported prediction $\hat{\mathbf{p}}_i$ is his expected likelihood for $\hat{\sigma}_j$. Agent i cannot pretend to have a different expected likelihood without reducing his expectation for his prediction score. Note that actually when agents play a *best prediction strategy profile* (See definition 2.13), each agent maximizes his individual prediction score given the signal strategy $(\theta_1, \theta_2, ..., \theta_n)$.

We cannot only pay agents based on prediction score, that will give no incentive to agents to report their private signals. Similar with BTS, we need "information score" which is motivated by a concept we call *Inconsistency*:

4.1 Inconsistency: Agree to Disagree

The common prior assumption tells us agents cannot agree to disagree. That is, if two agents receive the same private information, they must have the same "belief" about the world. In our setting, if agents tell the truth (or use a permutation strategy profile), whenever two agents report the same signal, they will report the same prediction as well. We use the concept of *inconsistency* to represent how much agents agree to disagree which we would like to discourage. We define *Inconsistency* as the "difference" between two random agents' predictions when they report the

same signal. Permutation strategy profiles will have the lowest inconsistency score 0. We will give a more formal definition later.

The *Inconsistency* concept motivates the *Truthful Mechanism*. We use a strictly proper scoring rule $PS(\cdot,\cdot)$ to define the "difference" between two agents' predictions in the *Truthful Mechanism* and give each agent i an **Information Score** by randomly picking an agent j and punishing agent i the "difference" between agent i's and agent j's predictions if they report the same signal.

Truthful Mechanism $\mathcal{M}$: Let $\alpha, \beta > 0$ be parameters and let PS be a strictly proper scoring rule, then we define the *truthful mechanism* $\mathcal{M}(\alpha, \beta, PS)$[1] as follows:

1. Each agent i reports a signal and a prediction $r_i = (\hat{\sigma}_i, \hat{\mathbf{p}}_i)$
2. For each agent i and agent j, we define a prediction score that depends on agent i's prediction and agent j's report signal
$$score_P(r_i, r_j) = PS(\hat{\sigma}_j, \hat{\mathbf{p}}_i),$$
and an information score
$$score_I(r_i, r_j) = \begin{cases} 0 & \hat{\sigma}_i \neq \hat{\sigma}_j \\ -(PS(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_j) - PS(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_i)) & \hat{\sigma}_i = \hat{\sigma}_j \end{cases}$$
3. Each agent i is matched with a random agent j. The payment for agent i is
$$payment_{\mathcal{M}(\alpha,\beta,\mathcal{PS})}(i, \mathbf{r}) = \alpha score_P(r_i, r_j) + \beta score_I(r_i, r_j).$$

Theorem 4.1. [2]

For any $\alpha, \beta > 0$ and any strictly proper scoring rule PS, $\mathcal{M}(\alpha, \beta, PS)$ has truth-telling as a strict Bayesian-Nash equilibrium whenever the prior Q is informative and symmetric.

Proof. We must show that for every agent, if other agents tell the truth, then this agent can (strictly) maximize his expected payoff if and only if he chooses to tell the truth.

Assume that all agents other than i are telling the truth. The probability that agent i is matched

[1] This mechanism is essentially the same as the Divergence-Based Bayesian Truth Serum mechanism previously published in [16].

[2] An equivalent theorem was proved in Radanovic and Faltings [16], and we include it here for completeness.

with agent j is $Pr(j) = \frac{1}{n-1}$. The expected payoff for agent i is:

$$\mathbb{E}[payment_{\mathcal{M}(\alpha,\beta,PS(\cdot,\cdot))}(i,\mathbf{r})|\sigma_i] \tag{12}$$

$$= \sum_{j\neq i}(Pr(j)\mathbb{E}[\alpha score_P(r_i,r_j) + \beta score_I(r_i,r_j)|\sigma_i]) \tag{13}$$

$$= \sum_{j\neq i}\frac{1}{n-1}[\alpha PS(\mathbb{E}(\hat{\sigma}_j|\sigma_i),\hat{\mathbf{p}}_i) + \beta(-Pr(\hat{\sigma}_j = \hat{\sigma}_i|\sigma_i)\mathbb{E}[(PS(\hat{\mathbf{p}}_j,\hat{\mathbf{p}}_j) - PS(\hat{\mathbf{p}}_j,\hat{\mathbf{p}}_i))|\sigma_i,\hat{\sigma}_j = \hat{\sigma}_i]] \tag{14}$$

$$= \sum_{j\neq i}\frac{1}{n-1}[\alpha PS(\mathbb{E}(\sigma_j|\sigma_i),\hat{\mathbf{p}}_i) + \beta(Pr(\sigma_j = \hat{\sigma}_i|\sigma_i)(PS(\mathbf{q}_{\hat{\sigma}_i},\hat{\mathbf{p}}_i) - PS(\mathbf{q}_{\hat{\sigma}_i},\mathbf{q}_{\hat{\sigma}_i})))] \tag{15}$$

$$= \alpha PS(\mathbb{E}(\sum_{j\neq i}\frac{1}{n-1}\sigma_j|\sigma_i),\hat{\mathbf{p}}_i) + \sum_{j\neq i}\frac{\beta}{n-1}[(Pr(\sigma_j = \hat{\sigma}_i|\sigma_i)(PS(\mathbf{q}_{\hat{\sigma}_i},\hat{\mathbf{p}}_i) - PS(\mathbf{q}_{\hat{\sigma}_i},\mathbf{q}_{\hat{\sigma}_i})))] \tag{16}$$

$$= \alpha PS(\theta_{-i}\mathbf{q}_{\sigma_i},\hat{\mathbf{p}}_i) + \sum_{j\neq i}\frac{\beta}{n-1}[(Pr(\sigma_j = \hat{\sigma}_i|\sigma_i)(PS(\mathbf{q}_{\hat{\sigma}_i},\hat{\mathbf{p}}_i) - PS(\mathbf{q}_{\hat{\sigma}_i},\mathbf{q}_{\hat{\sigma}_i})))] \tag{17}$$

$$= \alpha PS(\mathbf{q}_{\sigma_i},\hat{\mathbf{p}}_i) + \sum_{j\neq i}\frac{\beta}{n-1}[(Pr(\sigma_j = \hat{\sigma}_i|\sigma_i)(PS(\mathbf{q}_{\hat{\sigma}_i},\hat{\mathbf{p}}_i) - PS(\mathbf{q}_{\hat{\sigma}_i},\mathbf{q}_{\hat{\sigma}_i})))] \tag{18}$$

From (13) to (14): When $\hat{\sigma}_i \neq \hat{\sigma}_j$, the information score is 0, so we only need to consider the case $\hat{\sigma}_i = \hat{\sigma}_j$.

From (14) to (15): All agents other than i tell the truth, so $\hat{\sigma}_j = \sigma_j$ and

$$\hat{\mathbf{p}}_j = \mathbf{q}_{\sigma_j} = \mathbf{q}_{\hat{\sigma}_j} = \mathbf{q}_{\hat{\sigma}_i}.$$

From (15) to (16): The proper scoring rule is linear for the first entry.

From (16) to (17): Based on Claim 2.12, $E(\sum_{j\neq i}\frac{1}{n-1}\sigma_j|\sigma_i) = \theta_{-i}\mathbf{q}_{\sigma_i}$.

From (17) to (18): Note that for any $j \neq i$, agent j tells the truth so $\theta_{-i} = I$.

First, if agent i plays truthfully, then $\hat{\sigma}_i = \sigma_i, \hat{\mathbf{p}}_i = \mathbf{q}_{\sigma_i}$, and we will have $\mathbb{E}(payment(i,\mathcal{M})|\sigma_i) = \alpha PS(\mathbf{q}_{\sigma_i},\mathbf{q}_{\sigma_i})$ because $PS(\mathbf{q}_{\hat{\sigma}_i},\hat{\mathbf{p}}_i) - PS(\mathbf{q}_{\hat{\sigma}_i},\mathbf{q}_{\hat{\sigma}_i}) = 0$.

Now show that to receive a payment this high, agent i must play truthfully. Assume that $\mathbb{E}(payment(i,\mathcal{M})|\sigma_i) \geq \alpha PS(\mathbf{q}_{\sigma_i},\mathbf{q}_{\sigma_i})$. First, the second term of Equation (18) is non-positive based on the property of proper scoring rule. Then we must have that $PS(\mathbf{q}_{\sigma_i},\hat{\mathbf{p}}_i) \geq PS(\mathbf{q}_{\sigma_i},\mathbf{q}_{\sigma_i})$, but because PS is a strictly proper scoring rule, this happens only if $\hat{\mathbf{p}}_i = \mathbf{q}_{\sigma_i}$. But this implies that the second term of Equations (18) equals 0, and this requires that $PS(\mathbf{q}_{\hat{\sigma}_i},\hat{\mathbf{p}}_i) = PS(\mathbf{q}_{\hat{\sigma}_i},\mathbf{q}_{\hat{\sigma}_i})$.

However, by the properties of strictly proper scoring rules, this means $\mathbf{q}_{\hat{\sigma}_i} = \hat{\mathbf{p}}_i$. However, we already showed that $\hat{\mathbf{p}}_i = \mathbf{p}_i = \mathbf{q}_{\sigma_i}$. Putting this together we see that $\mathbf{q}_{\hat{\sigma}_i} = \mathbf{q}_{\sigma_i}$. Based on the informative prior assumption, this implies that $\hat{\sigma}_i = \sigma_i$.

So we proved that for any agent i, when other agents tell the truth, agent i can obtain the best expected payoff if and only if he tells the truth which means truth-telling is a strict Bayesian-Nash equilibrium in the truthful mechanism. □

4.2 Diversity from Disagreement

The previous section shows that we can encourage truth-telling (or permutation strategy profiles) by punishing *Inconsistency*. The problem is that there may be many other equilibria with inconsistency score 0. We consider an extreme example: All agents coordinate together and report the same signal σ. For every agent, he will predict the fraction of σ is 100% to maximize his prediction score. While such a strategy profile is consistent, the unitary predictions implies their report profiles have little information. Motivated by this extreme example, we define a concept called *Diversity*.

Recall that *Inconsistency* is the expected "difference" (we will describe this "difference" later) between two random agents' predictions when they report the same signal, we define *Diversity* as the expected *Hellinger-divergence* between two random agents' predictions when they report different signals. One of the reasons we use *Hellinger-divergence* is that it is f-divergence and we can use information monotonicity here. To give a flavor of our future proof, we will give a simple observation here which motivates our *Disagreement Mechanism*.

Observation 4.2. *If the number of agents is infinite and agents play a best prediction strategy, for every two agents i, j, their predictions becomes "closer" compared with agents play a permutation strategy profile.*

Proof. Based on Claim 2.12, agent i will report $\theta_{-i}\mathbf{q}_{\sigma_i}$ given σ_i is his private signal and recall that $\theta_{-i} = \frac{\sum_{j\neq i}\theta_i}{n-1}$ where $(\theta_1, \theta_2, ..., \theta_n)$ is the signal strategy. Since $\lim_{n\to\infty}\theta_{-i} = \lim_{n\to\infty}\theta_{-j}$, we can use θ to represent the common limit.

Via **information monotonicity** we have $D_f(\theta\mathbf{q}_{\sigma_i}, \theta\mathbf{q}_{\sigma_j}) \leq D_f(\mathbf{q}_{\sigma_i}, \mathbf{q}_{\sigma_j})$ which implies their reported predictions becomes "closer". □

Note that if the number of agents is not infinite, we can still have the same result if agents play a symmetric best prediction strategy profile since we only need $\theta_{-i} = \theta_{-j}$.

Ideally, we would like to show every equilibrium is a best prediction strategy profile, since it will imply any permutation strategy profile is more diverse than any other equilibrium. However, we believe it is not the case, instead, we introduce a new concept *Classification* that combines *Diversity* and *Inconsistency* and show that the permutation strategy profiles almost have the highest classification scores.

4.3 Classification: Diversity − Inconsistency

The set of all agents' report profiles is a classification. In general, a classification is a map $\mathcal{C} : U \to \Sigma$ where U are multiple points and Σ is a set of labels. All points that are mapped to one label σ construct a cluster C_σ. In our setting, the agents who report the same signal construct a cluster which is a collection of their reported profiles.

If all agents tell the truth, then the predictions that are labeled with the same label are in the same position in Δ_Σ (common prior assumption) while the predictions that are labeled with different labels are in different positions (because of our informative prior assumption). Intuitively, it is an ideal classification since typically we hope in a good classification, the points with the same label (in the same cluster) are close while the points with different labels are far from each other.

In general, a typical way to measure a classification $\mathcal{C}$ is

$$Q(\mathcal{C}) = \sum_{u,v\in U, \mathcal{C}(u)\neq\mathcal{C}(v)} D_1(u, v) - \sum_{u,v\in U, \mathcal{C}(u)=\mathcal{C}(v)} D_2(u, v)$$

where D_1 and D_2 are functions that measure the "difference" between points $u, v \in U$.

Note that for $Q(\mathcal{C}_r)$, $\sum_{u,v\in U,\mathcal{C}_r(u)\neq\mathcal{C}_r(v)} D_1(u,v)$ intuitively captures the concept of *Diversity* while $\sum_{u,v\in U,\mathcal{C}_r(u)\neq\mathcal{C}_r(v)} D_2(u,v)$ intuitively captures *Inconsistency*. So $Q(\mathcal{C}_r)$ captures the concept of *Diversity* $-$ *Inconsistency*.

We will also show that if we use *Diversity* $-$ *Inconsistency* as a new score, permutation strategies will have strictly higher score than any other symmetric equilibrium and if the number of agents are sufficient large, the equilibria with a higher score or even a score "close" to truth-telling must be "close" to a permutation strategy profile. We will use this property to design a mechanism that has truth-telling as both a robust-symmetric-quasi-focal and robust-approximate-quasi-focal equilibrium. We will design a mechanism that satisfies the below two conditions which implies that this mechanism has truth-telling as a weakly-quasi-focal equilibrium, and (1) has the same equilibria with our previous *Truthful Mechanism*; and (2) in which the average agent-welfare is *Diversity* $-$ *Inconsistency*.

We will see that the below *Disagreement Mechanism* we propose satisfies the above two conditions. To design this mechanism, we (a) first use a typical trick to create a zero-sum game which has the same equilibria with the previous *Truthful Mechanism*; (b) pay each agent an extra score that only depends on other agents which will not change the structure of the equilibria. We want this extra score to represent *Diversity* $-$ *Inconsistency*.

Disagreement Mechanism $\mathcal{M}+(\alpha,\beta,PS(\cdot,\cdot))$ $\mathbf{r}=\{r_1,r_2,...,r_n\}$ is all agents' report profiles where for any r, $r_i=(\hat{\sigma}_i,\hat{\mathbf{p}}_i)$.

1. *Zero-sum Trick*: Divide the agents into two non-empty groups-group A and group B. Each group of agents plays the game (mechanism) $\mathcal{M}$ that is restricted in their own group. For group A, each agent i_A receives a

$$\begin{aligned} score_{\mathcal{M}}(i_A,\mathbf{r}) =& payment_{\mathcal{M}(\alpha,\beta,PS(\cdot,\cdot))}(i_A,\mathbf{r}_A) \\ & -\frac{1}{|A|}\sum_{j_B\in B} payment_{\mathcal{M}(\alpha,\beta,PS(\cdot,\cdot))}(j_B,\mathbf{r}_B) \end{aligned}$$

   Where $payment_{\mathcal{M}(\alpha,\beta,PS(\cdot,\cdot))}(i_A,\mathbf{r}_A)$ is agent i_A's payment when he is paid by mechanism $\mathcal{M}(\alpha,\beta,PS(\cdot,\cdot))$ given group A's report profiles $\mathbf{r}_A$ and he can only be paired with a random peer from group A (we have similar explanation for $payment_{\mathcal{M}(\alpha,\beta,PS(\cdot,\cdot))}(j_B,\mathbf{r}_B)$). For agents in group B, we use the analogous way to score them.

2. *Additional Classification Reward*: Each agent i is matched with two random agents j,k chosen from all agents (including group A and group B), the payment for agent i is

$$payment_{\mathcal{M}+(\alpha,\beta,PS(\cdot,\cdot)}(i,\mathbf{r}) = score_{\mathcal{M}}(i,\mathbf{r}) + score_C(r_j,r_k)$$

   where

$$score_C(r_j,r_k)=\begin{cases} D^*(\hat{\mathbf{p}}_j,\hat{\mathbf{p}}_k) & \hat{\sigma}_j\neq\hat{\sigma}_k \\ -\sqrt{D^*(\hat{\mathbf{p}}_j,\hat{\mathbf{p}}_k)} & \hat{\sigma}_j=\hat{\sigma}_k \end{cases}$$

   recall that D^* denotes the Hellinger Divergence.

Note that our *Disagreement Mechanism* cannot make truth-telling *quasi-focal* since if there are m (the number of signals) agents and for any i, agent i always reports the i^{th} signal and predicts

that the j^{th} signal occurs with probability $\frac{1}{m-1}$ if $j \neq i$ while i^{th} signal occurs with probability 0. In this situation, each agent plays his best response and *Inconsistency* $= 0$, yet *Diversity* is much higher than it is in the truth-telling equilibrium. Instead, we are going to prove our *Disagreement Mechanism* makes truth-telling both a robustly-symmetric-quasi-focal and robustly-approximately-quasi-focal equilibrium.

Theorem 4.3 (Main Theorem)**.** *Given any SNIFE prior, if the number of agents $n \geq 3$, and if $\frac{\alpha}{\beta} < \frac{1}{4m}$ where m is the number of signals and n is the number of agents, then*

1. *$\mathcal{M}+(\alpha, \beta, PS(\cdot,\cdot))$ is truthful;*
2. *$\mathcal{M}+(\alpha, \beta, PS(\cdot,\cdot))$ has truth-telling as a* symmetric-quasi-focal *equilibrium;*
3. *$\mathcal{M}+(\alpha, \beta, PS(\cdot,\cdot))$ has truth-telling as a* (τ_1, γ_1)-robust-symmetric-quasi-focal *equilibrium; and*
4. *$\mathcal{M}+(\alpha, \beta, PS(\cdot,\cdot))$ has truth-telling as a* (τ_2, γ_2)-robust-approximate-quasi-focal *equilibrium*

where $\tau_1(\gamma_1) = O(\sqrt[3]{\gamma_1})$, $\gamma_2(n) = O(\frac{m}{\sqrt{n}})$ and $\tau_2(n) = O(\sqrt[6]{\frac{m^2}{n}})$ (the constants we omit only depend on the first two moments of prior Q)[3]*.*

Theorem 4.3 Part 1: $\mathcal{M}+(\alpha, \beta, PS(\cdot,\cdot))$ is truthful.

Claim 4.4. *The* Disagreement Mechanism *has the same equilibria as the* Truthful Mechanism*.*

Note that we have already show $\mathcal{M}(\alpha, \beta, PS(\cdot,\cdot))$ has truth-telling as a strict equilibrium for any SNIFE prior in Theorem 4.1. Since $\mathcal{M}+(\alpha, \beta, PS(\cdot,\cdot))$ does not change the equilibrium structure of $\mathcal{M}(\alpha, \beta, PS(\cdot,\cdot))$ according to Claim 4.4, we have $\mathcal{M}+(\alpha, \beta, PS(\cdot,\cdot))$ has truth-telling as a strict equilibrium for any SNIFE prior as well.

We will prove the rest of the theorem in the next section.

5 Showing Properties of Equilibrium

In this section, we are going to show Theorem 4.3 part 2, 3, 4.

We first give technical definitions for *Diversity* and *Inconsistency* and then prove that the average agent-welfare in the *Disagreement Mechanism* is *Diversity* $-$ *Inconsistency*.

We first introduce a short hand which will simplify the formula for *Diversity* and *Inconsistency*.

$$\int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k}) \triangleq \int_{\hat{\sigma}_j,\hat{\mathbf{p}}_j,\hat{\sigma}_k,\hat{\mathbf{p}}_k} Pr_{(\hat{\sigma}_j,\hat{\mathbf{p}}_j)\leftarrow s_j(\sigma_j)}(\hat{\sigma}_j,\hat{\mathbf{p}}_j) Pr_{(\hat{\sigma}_k,\hat{\mathbf{p}}_k)\leftarrow s_k(\sigma_k)}(\hat{\sigma}_k,\hat{\mathbf{p}}_k)$$

where s_j is the strategy of agent j and $s_j(\sigma_j)$ is a distribution over agent j's report profile $(\hat{\sigma}_j, \hat{\mathbf{p}}_j)$ given agent j receives private signal σ_j and uses strategy s_j, and similarly for agent k. This defines the natural measure on the reports of agents j and k given that they play strategies s_j and s_k and a fixed prior Q (which is implicit), and allows us to succinctly describe probabilities of events in this space.

[3]Actually $\tau_1(\gamma_1) = \frac{1}{c_1}\sqrt[3]{\frac{\gamma_1}{c_2,c_3,c_4}}$, $\gamma_2(n) = \frac{4\sqrt{2}m}{\sqrt{n}}$ and $\tau_2^6(n) = \frac{128*m^2}{nc_1^6(c_2c_3c_4)^2}$, $c_1 = \min_{s,t\in\Sigma} q(s|t)$, $c_2 = \min_{s,t\in\Sigma} Pr(s,t)$, $c_3 = \min_{u,v}\max_{s,t} ||\frac{q(u|s)}{q(u|t)} - \frac{q(v|s)}{q(v|t)}||^2$, $c_4 = \min_{s,t,u} f''(\frac{q(u|s)}{q(u|t)})$ where $f(x) = (\sqrt{x}-1)^2$.

Recall that *Diversity* is the expected Hellinger divergence D^* between two random agents when they report different signals, so

$$Diversity = \sum_{\substack{j \\ k \neq j}} \sum_{\sigma_j, \sigma_k} Pr(j,k) Pr(\sigma_j, \sigma_k) \int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k}) \delta(\hat{\sigma}_j \neq \hat{\sigma}_k) D^*(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_k)$$

where $Pr(j,k)$ is the probability agents j,k are picked, and $Pr(\sigma_j, \sigma_k)$ is the probability that agent j receives private signal σ_j and agent k receives private signal σ_k.

Similarly, we can write down the technical definition for *Inconsistency*. But here we do not use Hellinger divergence as the "difference" function in $\sum_{u,v \in U, \mathcal{C}_r(u)=\mathcal{C}_r(v)} D(u,v)$, we use square root of the Hellinger divergence which is the Hellinger distance as the "difference" function. The reason is we want to use the convexity of the Hellinger divergence and the triangle inequality of the Hellinger distance. We will describe the details in the future. For now we give a technical definition for *Inconsistency*:

$$Inconsistency = -\sum_{\substack{j \\ k \neq j}} \sum_{\sigma_j, \sigma_k} Pr(j,k) Pr(\sigma_j, \sigma_k) \int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k}) \delta(\hat{\sigma}_j = \hat{\sigma}_k) \sqrt{D^*(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_k)}$$

Now we define the *ClassificationScore* as the expected average extra score $score_C$:

$$ClassificationScore = \sum_{\substack{i \\ j \neq i}} \sum_{\substack{k \neq i,j \\ \sigma_i, \sigma_j, \sigma_k}} Pr(i) Pr(\sigma_i) Pr(j,k) Pr(\sigma_j, \sigma_k | \sigma_i) \int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k}) score_C(r_j, r_k)$$

Claim 5.1. *ClassificationScore = Diversity − Inconsistency*

Claim 5.2. *Any permutation strategy profile has the same ClassificationScore, Diversity, and Inconsistency as truth-telling.*

Claim 5.3. *The average agent-welfare in our* Disagreement Mechanism *is the ClassificationScore*

5.1 Proof Outline for Main Theorem

First note that the average agent-welfare is the *ClassificationScore*. We want to show that

1. If the number of agents is greater than 3, then any *symmetric* equilibrium that is not permutation equilibrium must have *ClassificationScore* strictly less than truth-telling; and any *symmetric* equilibrium that has *ClassificationScore* close to truth-telling must be close to a permutation equilibrium.

2. If the number of agents is sufficient large, then no equilibrium can have a *ClassificationScore* that is much greater than truth-telling; any equilibrium that has *ClassificationScore* close to truth-telling must be close to a permutation equilibrium.

To prove our main theorem, we first introduce the concept of *TotalDivergence* and then we use this value as a bridge. Recall that we defined

$Diversity = \sum_{j,k \neq j,\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k) \int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k})\delta(\hat{\sigma}_j \neq \hat{\sigma}_k)D^*(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_k)$, now we define a similar concept

$$TotalDivergence = \sum_{j,k,\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k) \int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k})D^*(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_k)$$

First note that total divergence is robust to summing over j, k or $j \neq k$ since when $j = k$, $D^*(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_k) = 0$.

We can see *TotalDivergence* $\geq$ *Diversity* since *TotalDivergence* also includes the divergence between the agents who report the same signals. We show that the equality holds if and only if $Inconsistency = 0$:

Claim 5.4. *For any strategy profile s, Diversity(s) = TotalDivergence(s) $\Leftrightarrow$ Inconsistency(s) = 0*

Corollary 5.5. *ClassificationScore(truth-telling)=Diversity(truth-telling)=TotalDivergence(truth-telling)*

Proof. At the truth-telling equilibrium, $\hat{\sigma}_i = \sigma_i, \hat{\mathbf{p}}_i = \mathbf{q}_{\sigma_i}$ for any i, so the inconsistency score of truth-telling is 0 since $\hat{\sigma}_j = \hat{\sigma}_k \Rightarrow \sigma_j = \sigma_k \Rightarrow \hat{\mathbf{p}}_j = \hat{\mathbf{p}}_k \Rightarrow D^*(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_k) = 0$ which implies this corollary. □

Now we begin to state our proof outline: For any equilibrium s, we define two modified strategies for s:

1. We define s_{BP} what we call a *best prediction strategy* of s as a strategy where each agent uses the same signal strategy which he uses in s but plays his *best prediction* which maximizes the prediction score. In this case, based on Claim 2.12, for any i, agent i plays $\theta_{-i}\mathbf{q}_{\sigma_i}$. We illustrate the report profiles of s_{BP} in the second picture from the left in Figure 1.

2. We define *symmetrized* s_{BP} as a strategy where each agent plays $\bar{\theta}_n\mathbf{q}_\sigma$ given σ is his private signal where $\bar{\theta}_n$ is the average signal strategy of s_{BP} (also of s). We show the report profiles of *symmetrized* s_{BP} in the third picture of Figure 1

Our proof can be divided in four parts which are illustrated in Figure 1:

(1) ***ClassificationScore***$(s) \leq$ ***TotalDivergence***(s_{BP})**.**[Lemma 5.6]. This is our main lemma and we defer the proof of main lemma to Section 5.3. Once we show it, we can directly prove that the modified mechanism has truth-telling as symmetric quasi-focal equilibrium based on information monotonicity which will be described in our proof for main theorem.

(2) ***TotalDivergence***$(s_{BP}) \approx$ ***TotalDivergence***$(symmetrized\ s_{BP})$ **when the number of agents is sufficient large.** [Lemma 5.7] Intuitively, when n is large enough, θ_{-j} will be close to $\bar{\theta}_n$. We will use this observation to prove this part.

(3) ***ClassificationScore***$(s^*) \geq$ ***ClassificationScore***$(truthtelling)$
$\Rightarrow$ ***TotalDivergence***$(truthtelling) \approx$ ***TotalDivergence***$(symmetrized\ s^*_{BP})$ **when the number of agents is sufficient large.**[Corollary 5.8] Here s^*_{BP} is the best prediction strategy of s^*. This part will also imply no equilibrium can have *ClassificationScore* that is much greater than truth-telling.

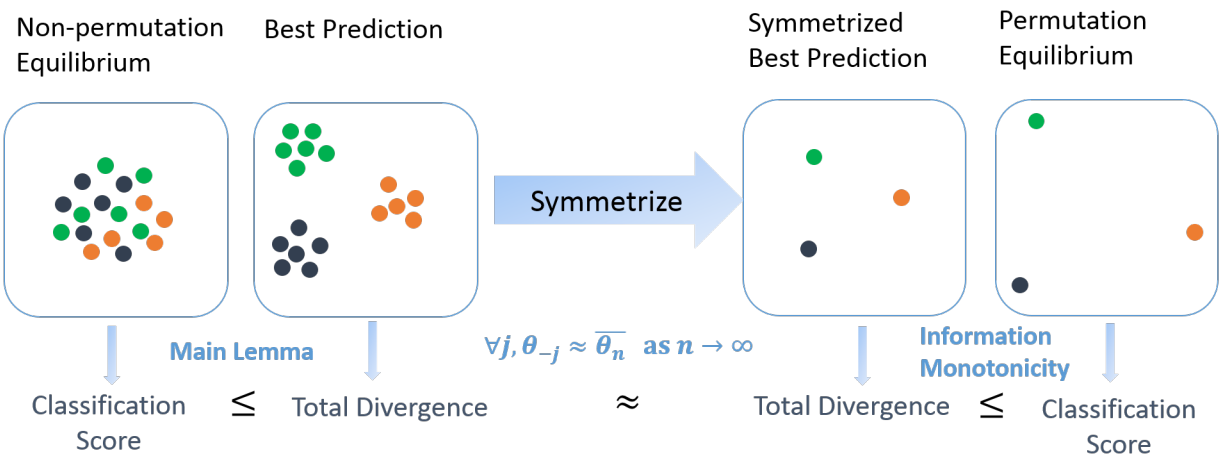

Figure 1: Proof Outline for Main Theorem

(4) ***TotalDivergence***$(truthtelling) \approx$ ***TotalDivergence***$(symmetrized\ s^*_{BP}) \Rightarrow \bar{\theta}_{\mathbf{n}} \approx \pi$ where $\bar{\theta}_{\mathbf{n}}$ is the average signal strategy of s^*. [Lemma 5.9]

We will show $TotalDivergence(truthtelling)$ is

$$\sum_{j,k,\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k)D^*(\mathbf{q}_{\sigma_j},\mathbf{q}_{\sigma_k})$$
$$= \sum_{j,k,\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k)D^*(\theta_\pi\mathbf{q}_{\sigma_j},\theta_\pi\mathbf{q}_{\sigma_k})$$

where $Pr(j,k)Pr(\sigma_j,\sigma_k)$ is the probability that agent j,k are picked and agent j receives private signal σ_j; agent k receives private signal σ_k.

We will also show $TotalDivergence(symmetrized\ s^*_{BP})$ is $\sum_{j,k,\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k)D^*(\bar{\theta}_n\mathbf{q}_{\sigma_j},\bar{\theta}_n\mathbf{q}_{\sigma_k})$.

Once we prove (1) and (2), we will show that (3) implies $\bar{\theta}_{\mathbf{n}} \approx \pi$. This, informally, means that if an equilibrium pays more than truth-telling, it must be close to a permutation equilibrium, and this pays about the same as truth-telling.

5.2 Proof for Main Theorem

(1) ***ClassificationScore***$(s) <$ ***TotalDivergence***(s_{BP})

Lemma 5.6 (Main Lemma)**.** *For any equilibrium s, if s_{BP} is a best prediction strategy of s, we have*

$$ClassificationScore(s) \leq TotalDivergence(s_{BP})$$

If the equality holds, then we have $Inconsistency(s) = 0$ and $s = s_{BP}$.

We defer the proof of our main lemma to Section 5.3.

Theorem 4.3 Part 2: $\mathcal{M}+(\alpha, \beta, PS(\cdot,\cdot))$ has truth-telling as a symmetric-quasi-focal equilibrium. We use our main lemma directly to prove: **any *symmetric* non-permutation equilibrium's agent welfare (*ClassificationScore*) must be strictly less than truth-telling**

Notice that if all agents play a symmetric signal strategy θ, then for any j, k, $\theta_{-j} = \theta_{-k} = \theta$. For any symmetric non-permutation equilibrium s, it is possible that the signal strategy of s is not a permutation or it is a permutation θ_π but agents do not report $\pi\mathbf{q}_\sigma$ given σ is their private signal. So we consider two cases:

(a) We first consider the case that the signal strategy θ of s is a permutation matrix θ_π, but agents do not report $\pi\mathbf{q}_\sigma$.

$$\begin{aligned}
\textit{ClassificationScore}(s) <& \textit{TotalDivergence}(s_{BP}) \\
=& \sum_{j,k,\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k)D^*(\theta_{-j}\mathbf{q}_{\sigma_j}, \theta_{-k}\mathbf{q}_{\sigma_k}) \\
=& \sum_{j,k,\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k)D^*(\theta_\pi\mathbf{q}_{\sigma_j}, \theta_\pi\mathbf{q}_{\sigma_k}) \\
=& \sum_{j,k,\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k)D^*(\mathbf{q}_{\sigma_j}, \mathbf{q}_{\sigma_k}) \\
=& \textit{TotalDivergence}(truthtelling) \\
=& \textit{ClassificationScore}(truthtelling)
\end{aligned}$$

The first inequality follows from our main lemma. The inequality is strict for the following reason: when the signal strategy θ of s is a permutation matrix, s_{BP} is a permutation strategy profile since for any i, agent i's best prediction is $\theta_{-i}\mathbf{q}_\sigma = \theta\mathbf{q}_\sigma$. Based on our main lemma if $\textit{ClassificationScore}(s) = \textit{TotalDivergence}(s_{BP})$, we have $s = s_{BP}$ which implies that s is a permutation strategy profile which is a contradiction to the fact s is a non-permutation strategy profile.

The second line follows since at s_{BP}, each agent's reported prediction only depends on his private signal.

The last equality follows from Corollary 5.5.

(b) We consider the case that the signal strategy θ of s is not a permutation matrix. The above proof still holds except in two places: one is the inequality in the first line may not be strict; another is the equality in the fourth line should be a strict inequality:

$$\begin{aligned}
&\sum_{j,k,\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k)D^*(\theta_\pi\mathbf{q}_{\sigma_j}, \theta_\pi\mathbf{q}_{\sigma_k}) \\
<& \sum_{j,k,\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k)D^*(\mathbf{q}_{\sigma_j}, \mathbf{q}_{\sigma_k})
\end{aligned}$$

The inequality must be strict since based on Corollary 2.19, we know that if θ is not a permutation, and Q is fine-grained, then there exists $\sigma_1 \neq \sigma_2$ such that $D^*(\theta\mathbf{q}_{\sigma_1}, \theta\mathbf{q}_{\sigma_2}) < D^*(\mathbf{q}_{\sigma_1}, \mathbf{q}_{\sigma_2})$. Also based on non-zero assumption of Q, we have $Pr(\sigma_j = \sigma_1, \sigma_k = \sigma_2) > 0$.

So in both of the above two cases, we have

$$\textit{ClassificationScore}(s) < \textit{ClassificationScore}(truthtelling)$$

if s is not a permutation equilibrium

Things are more complicated if agents play asymmetric equilibria, and we begin to handle this case by showing Lemma 5.7, Corollary 5.8, and Lemma 5.9.

(2) *TotalDivergence*(s_{BP}) $\approx$ *TotalDivergence*($symmetrized\ s_{BP}$) when the number of agents is sufficient large. We symmetrize s_{BP} which means we let each agent i report $\bar{\theta}_n \mathbf{q}_{\sigma_i}$ given σ_i is agent i's private signal and $\bar{\theta}_n$ is the average signal strategy of s_{BP} and show that the total divergence will not change much. Intuitively, this is because θ_{-i} are similar among agents when there are many agents.

Lemma 5.7. *Given any SNIFE prior Q, for any $\epsilon > 0$, there exists $N_\epsilon = \frac{32*m^2}{\epsilon^2}$ such that if $n > N_\epsilon$, for any strategy $(\theta_1, \theta_2, ..., \theta_n)$, any two agents j, k,*

$$|D^*(\theta_{-j}\mathbf{q}_{\sigma_j}, \theta_{-k}\mathbf{q}_{\sigma_k}) - D^*(\bar{\theta}_n\mathbf{q}_{\sigma_j}, \bar{\theta}_n\mathbf{q}_{\sigma_k})| < \epsilon$$

Proof of Lemma 5.7. For convenience, let $s = \sigma_j, t = \sigma_k$

$$|D^*(\theta_{-j}\mathbf{q}_s, \theta_{-k}\mathbf{q}_t) - D^*(\bar{\theta}_n\mathbf{q}_s, \bar{\theta}_n\mathbf{q}_t)| \tag{19}$$

$$=|\sum_u \left(\sqrt{\theta_{-j}(u,\cdot)\mathbf{q}_s} - \sqrt{\theta_{-k}(u,\cdot)\mathbf{q}_t}\right)^2 - \left(\sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_s} - \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_t}\right)^2| \tag{20}$$

$$=|\sum_u \left(\sqrt{\theta_{-j}(u,\cdot)\mathbf{q}_s} - \sqrt{\theta_{-k}(u,\cdot)\mathbf{q}_t} - \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_s} + \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_t}\right) *$$
$$\left(\sqrt{\theta_{-j}(u,\cdot)\mathbf{q}_s} - \sqrt{\theta_{-k}(u,\cdot)\mathbf{q}_t} + \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_s} - \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_t}\right)| \tag{21}$$

$$\leq 2 * m * \max_u \left(|\sqrt{\theta_{-j}(u,\cdot)\mathbf{q}_s} - \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_s}| + |\sqrt{\theta_{-k}(u,\cdot)\mathbf{q}_t} - \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_t}|\right) \tag{22}$$

$$\leq 4 * m * \max_{u,s,j} |\sqrt{\theta_{-j}(u,\cdot)\mathbf{q}_s} - \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_s}| \tag{23}$$

The first equality follows from the definition of Helinger-divergence.

The second equality is just formula for the difference of square.

To arrive at (22), $\sum_u |\left(\sqrt{\theta_{-j}(u,\cdot)\mathbf{q}_s} - \sqrt{\theta_{-k}(u,\cdot)\mathbf{q}_t} + \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_s} - \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_t}\right)| \leq \sum_u 2 = 2m$ where the inequality follows from the fact $0 < D^* < 1$.

The last equality follows since both $|\sqrt{\theta_{-j}(u,\cdot)\mathbf{q}_s} - \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_s}|$ and $|\sqrt{\theta_{-k}(u,\cdot)\mathbf{q}_t} - \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_t}|$ are less than $\max_{u,s,j} |\sqrt{\theta_{-j}(u,\cdot)\mathbf{q}_s} - \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_s}|$.

Now we consider two cases for any u, s, j:

(1) $|\sqrt{\theta_{-j}(u,\cdot)\mathbf{q}_s} - \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_s}| \leq \frac{\epsilon}{4*m}$: It is clear the result in this Lemma follows.

(2) $|\sqrt{\theta_{-j}(u,\cdot)\mathbf{q}_s} - \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_s}| > \frac{\epsilon}{4*m}$: Notice that $(n-1)\theta_{-j} = n\bar{\theta}_n - \theta_j$, then we can see

$$\theta_{-j} = \bar{\theta}_n + \frac{1}{n}(\theta_{-j} - \theta_j)$$

$$4 * m * |\sqrt{\theta_{-j}(u,\cdot)\mathbf{q}_s} - \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_s}| \tag{24}$$

$$= 4 * m * |\frac{\theta_{-j}(u,\cdot)\mathbf{q}_s - \bar{\theta}_n(u,\cdot)\mathbf{q}_s}{\sqrt{\theta_{-j}(u,\cdot)\mathbf{q}_s} + \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_s}}| \tag{25}$$

$$= 4 * m * \frac{\frac{1}{n}|(\theta_{-j}(u,\cdot) - \theta_j(u,\cdot))\mathbf{q}_s|}{\sqrt{\theta_{-j}(u,\cdot)\mathbf{q}_s} + \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_s}} \tag{26}$$

$$< 4 * m * 2 * \frac{4 * m}{\epsilon}\frac{1}{n} < \epsilon \tag{27}$$

when $n > N_\epsilon = \frac{32*m^2}{\epsilon^2}$

The first equality follows from the formula of the difference of squares.

The second equality follows from $\theta_{-j} = \bar{\theta}_n + \frac{1}{n}(\theta_{-j} - \theta_j)$.

If $|\sqrt{\theta_{-j}(u,\cdot)\mathbf{q}_s} - \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_s}| > \frac{\epsilon}{4*m}$, we have $|\sqrt{\theta_{-j}(u,\cdot)\mathbf{q}_s} + \sqrt{\bar{\theta}_n(u,\cdot)\mathbf{q}_s}| > \frac{\epsilon}{4*m}$ as well, the third line follows.

□

(3) *ClassificationScore*$(s^*) \geq$ *ClassificationScore*$(truthtelling)$
$\Rightarrow$ *TotalDivergence*$(truthtelling) \approx$ *TotalDivergence*$(symmetrized\ s^*_{BP})$ when the number of agents is sufficient large. The below corollary is derived from Lemma 5.7. It will imply not only
TotalDivergence$(truthtelling) \approx$ *TotalDivergence*$(symmetrized\ s^*_{BP})$ but also any equilibrium cannot have agent-welfare (*ClassificationScore*) that is much greater than truth-telling when the number of agents is sufficient large.

Corollary 5.8. *Given any SNIFE prior Q, for any $\epsilon > 0$, if $n > N_\epsilon = \frac{128*m^2}{\epsilon^2}$, for any equilibrium s^* that has greater ClassificationScore than the truth-telling ClassificationScore minus $\epsilon/2$:*

$$\begin{aligned}
&\textit{Classification}(truthtelling) \\
&< \textit{Classification}(s^*) + \frac{\epsilon}{2} \\
&< \textit{TotalDivergence}(symmetrized\ s^*_{BP}) + \epsilon \\
&\leq \textit{Classification}(truthtelling) + \epsilon
\end{aligned}$$

Proof for Corollary 5.8.

$$\begin{aligned}
\textit{TotalDivergence}(truthtelling) &= \textit{ClassificationScore}(truthtelling) \\
&\leq \textit{ClassificationScore}(s^*) + \frac{\epsilon}{2} \\
&\leq \textit{TotalDivergence}(s^*_{BP}) + \frac{\epsilon}{2} \\
&< \textit{TotalDivergence}(symmetrized\ s^*_{BP}) + \epsilon \\
&\leq \textit{ClassificationScore}(truthtelling) + \epsilon
\end{aligned}$$

The first equality follows from Corollary 5.5.
The second inequality follows from the condition.
The third inequality follows from the main lemma.
The fourth inequality follows from Lemma 5.7.
The last inequality follows from information monotonicity since

$$\begin{aligned}
&\textit{ClassificationScore}(truthtelling) - \textit{TotalDivergence}(symmetrized\ s^*_{BP}) \\
&= \sum_{\substack{j \\ k \neq j}} \sum_{\sigma_j, \sigma_k} Pr(j,k) Pr(\sigma_j, \sigma_k) D^*(\mathbf{q}_{\sigma_j}, \mathbf{q}_{\sigma_k}) - \sum_{j,k,\sigma_j,\sigma_k} Pr(j,k) Pr(\sigma_j, \sigma_k) D^*(\bar{\theta}_n \mathbf{q}_{\sigma_j}, \bar{\theta}_n \mathbf{q}_{\sigma_k}) \\
&= \sum_{j,k,\sigma_j,\sigma_k} Pr(j,k) Pr(\sigma_j, \sigma_k) (D^*(\mathbf{q}_{\sigma_j}, \mathbf{q}_{\sigma_k}) - D^*(\bar{\theta}_n \mathbf{q}_{\sigma_j}, \bar{\theta}_n \mathbf{q}_{\sigma_k})) \geq 0
\end{aligned}$$

The second equality follows since if $j = k$, $D^*(\mathbf{q}_{\sigma_j}, \mathbf{q}_{\sigma_k}) = 0$

□

This corollary induces the following result:

No equilibrium can have agent-welfare that is much greater than truth-telling Let $\frac{\epsilon}{2} = \gamma_2$, we need $n \geq \frac{128*m^2}{\epsilon^2}$ to obtain γ_2 tolerance based on Corollary 5.8. By manipulations, we obtain our result.

(3) $\Rightarrow$ (4) $\bar{\theta}_{\mathbf{n}} \approx \pi$ We already know that if s^* obtains higher *ClassificationScore* than truth-telling, the classification score of truth-telling is close to that of the symmetrized s^*_{BP}. Now we will prove that $\bar{\theta}_n$ is close to a permutation where $\bar{\theta}_n$ is the average signal strategy of s^*. We prove it by contradiction. We first assume that $\bar{\theta}_n$ is far from a permutation equilibrium, that is, recalling the definition of τ-close, we assume there exists a row of $\bar{\theta}_n$ that has at least two large numbers. Formally, we assume that given any τ, there exists $u', v', w' \in \Sigma$ such that $\bar{\theta}_n(u', v') > \tau, \bar{\theta}_n(u', w') > \tau$. We will prove that it is impossible when $n > N(\tau, Q)$ since when $\bar{\theta}_n$ is far from a permutation, the total divergence of symmetrized s_{BP} is far from the classification score of truth-telling which contradicts Corollary 5.8.

The below lemma tells us if a symmetric strategy s_θ, where agents play θ as their signal strategy and best prediction as their reported prediction, is far from a permutation, then the total divergence of s_θ and total divergence of permutation are also far. Once we proved the below lemma, we can replace θ by $\bar{\theta}_n$ to finish our proof for Theorem 4.3 part 4, that is, truth-telling is robust-approximate-quasi-focal.

Lemma 5.9. *Given any fixed τ, for any signal strategy θ, if there exists $u', v', w' \in \Sigma$ such that $\theta(u', v') > \tau, \theta(u', w') > \tau$, then consider the case s_θ that all agents play θ as their signal strategy and report their best prediction, we have*

$$\textit{TotalDivergence}(\textit{truthtelling}) - \textit{TotalDivergence}(s_\theta) \geq c_2(\tau c_1)^3 c_4 c_3$$

Proof of Lemma 5.9. We first write *TotalDivergence* in an explicit form:

$$\sum_{j,k,\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k)D^*(\mathbf{q}_{\sigma_j}, \mathbf{q}_{\sigma_k}) - \sum_{j,k,\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k)D^*(\theta\mathbf{q}_{\sigma_j}, \theta\mathbf{q}_{\sigma_k}) \tag{28}$$

Actually, We will show for any j, k,

$$\sum_{\sigma_j,\sigma_k} Pr(\sigma_j,\sigma_k)D^*(\mathbf{q}_{\sigma_j}, \mathbf{q}_{\sigma_k}) - \sum_{\sigma_j,\sigma_k} Pr(\sigma_j,\sigma_k)D^*(\theta\mathbf{q}_{\sigma_j}, \theta\mathbf{q}_{\sigma_k}) \tag{29}$$

is greater than $c_2(\tau c_1)^3 c_4 c_3$, which implies the result.

We want give a lower bound for (29). In order to obtain this lower bound, we are going to transform this value to $\sum_u \lambda_u g(x_u) - g(\sum_u \lambda_u x_u)$ where $g(\cdot)$ is a convex function. To obtain a lower bound of $\sum_u \lambda_u g(x_u) - g(\sum_u \lambda_u x_u)$, we have an observation:

For any convex function $g(\cdot)$, $g(\sum_u \lambda_u x_u)$ and $\sum_u \lambda_u g(x_u)$ are "very different" if there are two large coefficients λ_1 and λ_2 with the corresponding x_1 and x_2 that are "very different". Now we introduce a claim to show this observation.

Claim 5.10.

$$\sum_u \lambda_u g(x_u) - g(\sum_u \lambda_u x_u) \geq \frac{d_2(g)}{2} \frac{\lambda_1\lambda_2}{\lambda_1+\lambda_2} ||x_1 - x_2||^2$$

where $d_2(g)$ is a lower bound of $g''(\cdot)$

Proof.

$$g\left(\sum_u \lambda_u x_u\right) \leq (\lambda_1+\lambda_2) g\left(\frac{\lambda_1 x_1+\lambda_2 x_2}{\lambda_1+\lambda_2}\right) + \sum_{u>2} \lambda_u g(x_u) \leq \sum_u \lambda_u g(x_u)$$

So

$$\begin{aligned}
&\sum_u \lambda_u g(x_u) - g(\sum_u \lambda_u x_u) \\
&\geq \sum_u \lambda_u g(x_u) - (\lambda_1+\lambda_2) g\left(\frac{\lambda_1 x_1+\lambda_2 x_2}{\lambda_1+\lambda_2}\right) - \sum_{u>2} \lambda_u g(x_u) \\
&= \lambda_1 g(x_1) + \lambda_2 g(x_2) - (\lambda_1+\lambda_2) g(\frac{\lambda_1 x_1+\lambda_2 x_2}{\lambda_1+\lambda_2}) \\
&= (\lambda_1+\lambda_2)(\frac{\lambda_1 g(x_1)+\lambda_2 g(x_2)}{\lambda_1+\lambda_2} - g(\frac{\lambda_1 x_1+\lambda_2 x_2}{\lambda_1+\lambda_2})) \\
&\geq (\lambda_1+\lambda_2)\frac{d_2(g)}{2}\frac{\lambda_1\lambda_2}{(\lambda_1+\lambda_2)^2}||x_1-x_2||^2 \\
&= \frac{d_2(g)}{2}\frac{\lambda_1\lambda_2}{\lambda_1+\lambda_2}||x_1-x_2||^2
\end{aligned}$$

where $d_2(g)$ is the lower bound of $g''(\cdot)$

The first inequality follows if we rewrite $\sum_u \lambda_u x_u$ as $(\lambda_1+\lambda_2)\frac{\lambda_1 x_1+\lambda_2 x_2}{\lambda_1+\lambda_2} + \sum_{u>2} \lambda_u x_u$ and apply convexity.

Then we do several manipulations including taking $\lambda_1+\lambda_2$ outside. For continuous convex function $g(\cdot)$, we have $tg(x)+(1-t)g(y)-g(tx+(1-t)y) \geq \frac{1}{2}d_2(g)t(1-t)||x-y||^2$ according to [13], then we replace t by $\frac{\lambda_1}{\lambda_1+\lambda_2}$ and set $x=x_1, y=x_2$ and obtain the final result.

□

We can think of $\theta(u',v')$ and $\theta(u',w')$ as the two large coefficients (actually they are part of the coefficients). Then we need to find two "very different" entries that corresponding to those large coefficients. We pick two specific signals $s',t' \in \Sigma$ such that $\mathbf{q}_{s'}$ and $\mathbf{q}_{t'}$ are "very different" in position v' and w'. The reason we do this is that when we compute $\theta\mathbf{q}$, $\theta(u',v')$ and $\theta(u',w')$ are the two large entries which correspond to the positions v' and w' in $\mathbf{q}$. Formally, we pick $s',t' \in \Sigma$ such that

$$\left\|\frac{q(v'|s')}{q(v'|t')} - \frac{q(w'|s')}{q(w'|t')}\right\| = \max_{s,t}\left\|\frac{q(v'|s)}{q(v'|t)} - \frac{q(w'|s)}{q(w'|t)}\right\|$$

Once we have chosen the two specific signals, since $Pr(s',t')(D^*(\mathbf{q}_{s'},\mathbf{q}_{t'}) - D^*(\theta\mathbf{q}_{s'},\theta\mathbf{q}_{t'}))$ is less than (29) based on the fact $D^*(\mathbf{q}_s,\mathbf{q}_t) - D^*(\theta\mathbf{q}_s,\theta\mathbf{q}_t) \geq 0$ for $s,t \neq s',t'$, we will give a lower bound of $Pr(s',t')(D^*(\mathbf{q}_{s'},\mathbf{q}_{t'}) - D^*(\theta\mathbf{q}_{s'},\theta\mathbf{q}_{t'}))$ which is also a lower bound of (29).

Let $f(x) = (\sqrt{x}-1)^2$. For convenience, we will write the dot product of two vectors $\sum_v a(v)b(v)$ as $a(\cdot)b(\cdot)$. Now we give a explicit form of D^*:

$$Pr(s',t')(D^*(\mathbf{q}_{s'},\mathbf{q}_{t'}) - D^*(\theta\mathbf{q}_{s'},\theta\mathbf{q}_{t'})) \tag{30}$$

$$= Pr(s',t')\left(\sum_v q(v|s') f\left(\frac{q(v|t')}{q(v|s')}\right) - \sum_u \theta(u,\cdot)q(\cdot|s') f\left(\frac{1}{\theta(u,\cdot)q(\cdot|s')}\theta(u,\cdot)q(\cdot|t')\right)\right) \tag{31}$$

We take $\sum_u \theta(u,\cdot)q(\cdot|s')$ out and note that $\sum_u \theta(u,v)=1$, so $\sum_u \theta(u,\cdot)q(\cdot|s')\frac{1}{\theta(u,\cdot)q(\cdot|s')}\theta(u,v)=1$, then we obtain (32) from (31).

$$(31) = Pr(s',t')\sum_u \theta(u,\cdot)q(\cdot|s')* \left(\frac{1}{\theta(u,\cdot)q(\cdot|s')}\sum_v \theta(u,v)q(v|s')f\left(\frac{q(v|t')}{q(v|s')}\right) - f\left(\frac{1}{\theta(u,\cdot)q(\cdot|s')}\sum_v \theta(u,v)q(v|s')\frac{q(v|t')}{q(v|s')}\right)\right) \tag{32}$$

Then we pick the special u' to obtain (33). For the part $\sum_{u\neq u'}$, since $f(\cdot)$ is a convex function, we have

$$\frac{1}{\theta(u,\cdot)q(\cdot|s')}\sum_v \theta(u,v)q(v|s')f(\frac{q(v|t')}{q(v|s')}) \geq f\left(\frac{1}{\theta(u,\cdot)q(\cdot|s')}\sum_v \theta(u,v)q(v|s')\frac{q(v|t')}{q(v|s')}\right)$$

so (32) is greater than (33).

$$(32) \geq Pr(s',t')\theta(u',\cdot)q(\cdot|s')* \left(\frac{1}{\theta(u',\cdot)q(\cdot|s')}\sum_v \theta(u',v)q(v|s')f\left(\frac{q(v|t')}{q(v|s')}\right) - f\left(\frac{1}{\theta(u',\cdot)q(\cdot|s')}\sum_v \theta(u',v)q(v|s')\frac{q(v|t')}{q(v|s')}\right)\right) \tag{33}$$

Note that $\theta(u',v')$ and $\theta(u',w')$ are large, so in the convex function $f(\cdot)$, there are two large coefficients $\frac{1}{\theta(u',\cdot)q(\cdot|s')}\theta(u',v')q(v'|s')$ and $\frac{1}{\theta(u',\cdot)q(\cdot|s')}\theta(u',w')q(w'|s')$ which correspond to $\frac{q(v'|t')}{q(v'|s')}$ and $\frac{q(w'|t')}{q(w'|s')}$. Then based on our choice for s',t' and Claim 5.10, we have

$$(33) \geq Pr(s',t')\theta(u',\cdot)q(\cdot|s')\frac{c_4}{2}\left(\frac{(\theta(u',v')q(v'|s'))*(\theta(v',w')q(w'|s'))}{\theta(u',v')q(v'|s')+\theta(v',w')q(w'|s')}\left\|\frac{q(v'|t')}{q(v'|s')}-\frac{q(w'|t')}{q(w'|s')}\right\|^2\right) \tag{34}$$

$$\geq c_2(\tau c_1)^3 c_4 c_3 \tag{35}$$

The last inequality follows since $Pr(s',t')\geq c_2$, both $\theta(u',v')q(v'|s')$ and $\theta(v',w')q(w'|s')$ are greater than τc_1. Also note that:

$$\theta(u',\cdot)q(\cdot|s') \geq \theta(u',v')q(v'|s')+\theta(v',w')q(w'|s') \geq 2\tau c_1$$

and

$$\theta(u',v')q(v'|s')+\theta(v',w')q(w'|s') \leq 1.$$

□

Theorem 4.3 Part 3: $\mathcal{M}+(\alpha,\beta,PS(\cdot,\cdot))$ **has truth-telling as a** ***(τ_1,γ_1)-robust-symmetric-quasi-focal*****:**

Any *symmetric* equilibrium that has agent-welfare close to truth-telling must be close to a permutation equilibrium: We have already proved that no symmetric equilibrium pays more than truth-telling. For the symmetric equilibrium s^* such that $\mathit{ClassificationScore}(s^*) > \mathit{ClassificationScore}(\mathit{truthtelling}) - \gamma_1$, we have

$$\mathit{TotalDivergence}(\mathit{truthtelling}) = \mathit{ClassificationScore}(\mathit{truthtelling}) \leq \mathit{ClassificationScore}(s^*) + \gamma_1$$
$$\leq \mathit{TotalDivergence}(s^*_{BP}) + \gamma_1 \leq \mathit{ClassificationScore}(\mathit{truthtelling}) + \gamma_1$$

Let $\gamma_1 = (\tau_1 c_1)^3 c_2 c_3 c_4$, then s^* is τ_1 close to a permutation equilibrium or there will be a contradiction based on Lemma 5.9. By manipulations, we will obtain our result.

Theorem 4.3 Part 4: $\mathcal{M}+(\alpha, \beta, PS(\cdot,\cdot))$ has truth-telling as a *(τ_2, γ_2)-robust-approximate-quasi-focal*:

If the number of agents is sufficient large, any equilibrium that has agent-welfare close to truth-telling must be close to permutation equilibrium: Let $\epsilon = (\tau_2 c_1)^3 c_2 c_3 c_4$, if $n > \frac{32 * m^2}{(\epsilon/2)^2}$, we have already proved that $\mathit{TotalDivergence}(\mathit{truthtelling}) - \mathit{TotalDivergence}(\mathit{symmetrized}\ s^*_{BP}) < \epsilon$ based on Corollary 5.8. If s^* is not τ_2 close to a permutation equilibrium, we will have

$$\mathit{TotalDivergence}(\mathit{truthtelling}) - \mathit{TotalDivergence}(\mathit{symmetrized}\ s^*_{BP}) > (\tau_2 c_1)^3 c_2 c_3 c_4 = \epsilon$$

which is a contradiction based on Lemma 5.9. By manipulations, we obtain our result.

5.3 Proof for Main Lemma

In this section, we will pove the main lemma: the classification score of non-permutation equilibrium s is less than the total divergence of the report profiles when agents report their best predictions given they still use the signal strategy of s. We first show the inequality and then show that if the equality holds, then s is consistent and $s = s_{BP}$.

In order to show the inequality, we first show

$$\mathit{TotalDivergence}(s) - \mathit{TotalDivergence}(s_{BP}) \leq \mathit{Inconsistency}(s)$$

once we show this, since we have $\mathit{ClassificationScore} = \mathit{Diversity} - \mathit{Inconsistency}$ and $\mathit{Diversity} \leq \mathit{TotalDivergence}$, our main lemma $\mathit{ClassificationScore}(s) \leq \mathit{TotalDivergence}(s_{BP})$ will follow since

$$\begin{aligned}\mathit{ClassificationScore}(s) =& \mathit{Diversity}(s) - \mathit{Inconsistency}(s) \\ \leq& \mathit{TotalDivergence}(s) - \mathit{Inconsistency}(s) \leq \mathit{TotalDivergence}(s_{BP}) \end{aligned} \quad (36)$$

To prove $\mathit{TotalDivergence}(s) - \mathit{TotalDivergence}(s_{BP}) \leq \mathit{Inconsistency}(s)$, we will write it in a explicit form:

$$\mathit{TotalDivergence}(s) - \mathit{TotalDivergence}(s_{BP}) \quad (37)$$
$$= \sum_{j,k,\sigma_j,\sigma_k} Pr(j,k) Pr(\sigma_j, \sigma_k) \int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k}) (D^*(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_k) - D^*(\theta_{-j}\mathbf{q}_{\sigma_j}, \theta_{-k}\mathbf{q}_{\sigma_k})) \quad (38)$$

It is difficult to compare $D^*(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_k)$ and $D^*(\theta_{-j}\mathbf{q}_{\sigma_j}, \theta_{-k}\mathbf{q}_{\sigma_k})$ directly. To deal with this problem, we introduce a new value $D^*(\hat{\mathbf{p}}_j, \theta_{-k}\mathbf{q}_{\sigma_k})$ and write (37) as

$$\sum_{j,k,\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k)*$$
$$\int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k}) \left(D^*(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_k) - D^*(\hat{\mathbf{p}}_j, \theta_{-k}\mathbf{q}_{\sigma_k}) + D^*(\hat{\mathbf{p}}_j, \theta_{-k}\mathbf{q}_{\sigma_k}) - D^*(\theta_{-j}\mathbf{q}_{\sigma_j}, \theta_{-k}\mathbf{q}_{\sigma_k})\right) \tag{39}$$

We will first give the analysis for $D^*(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_k) - D^*(\hat{\mathbf{p}}_j, \theta_{-k}\mathbf{q}_{\sigma_k})$, then we will see $D^*(\hat{\mathbf{p}}_j, \theta_{-k}\mathbf{q}_{\sigma_k}) - D^*(\theta_{-j}\mathbf{q}_{\sigma_j}, \theta_{-k}\mathbf{q}_{\sigma_k})$ is similar.

Remember that both $D^*(a, \cdot)$ and $D^*(\cdot, b)$ are convex functions. So $D^*(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_k) - D^*(\hat{\mathbf{p}}_j, \theta_{-k}\mathbf{q}_{\sigma_k})$ can be seen as $g(\hat{\mathbf{p}}_k) - g(\theta_{-k}\mathbf{q}_{\sigma_k})$ where $g(\cdot)$ is convex function $D^*(\hat{\mathbf{p}}_j, \cdot)$.

Recall that

$$Inconsistency = \sum_{\substack{j \\ k \neq j}} \sum_{\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k) \int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k})\delta(\hat{\sigma}_j = \hat{\sigma}_k)\sqrt{D^*(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_k)}$$

We hope we can obtain a upper bound for $g(\hat{\mathbf{p}}_k) - g(\theta_{-k}\mathbf{q}_{\sigma_k})$ that relates to agent k's neighbors' best response predictions. Here agent k's *neighbors* mean the agents who report the same signal with agent k and *best response prediction* means the reported prediction at equilibrium.

Now we begin to analyze the relationship between $\hat{\mathbf{p}}_k$ and $\theta_{-k}\mathbf{q}_{\sigma_k}$. Recall that each agent's payment depends on his prediction score and information score. $\theta_{-k}\mathbf{q}_{\sigma_k}$ maximizes the prediction score while $\hat{\mathbf{p}}_k$ maximizes the payment. The information score depends on agent k's neighbors' reported predictions $\{\hat{\mathbf{p}}_l | l \neq k\}$. So we can see $\hat{\mathbf{p}}_k$ is related to both his best prediction $\theta_{-k}\mathbf{q}_{\sigma_k}$ and his neighbors' reported predictions $\{\hat{\mathbf{p}}_l | l \neq k\}$. Actually we will show that $\hat{\mathbf{p}}_k$ can be computed as a linear combination of $\theta_{-k}\mathbf{q}_{\sigma_k}$ and $\{\hat{\mathbf{p}}_l | l \neq k\}$, which is based on the fact that every proper scoring rule is linear for the first entry (we will discuss the detail in the below proof). Once we have this result, we can construct a linear system about agents' reported predictions $\{\hat{\mathbf{p}}_i | i\}$ and their best predictions. This linear system helps us obtain a upper bound for $g(\hat{\mathbf{p}}_k) - g(\theta_{-k}\mathbf{q}_{\sigma_k})$ which upper-bounds the distance between agent k's best response prediction and his neighbors' best response predictions.

Equilibrium Analysis We will analyze the equilibrium in our *Truthful Mechanism* which is also the equilibrium in our *Disagreement Mechanism*. We first show, in Claim 5.11, that at equilibrium, an agent's reported prediction only depends on his private signal and reported signal. Then we use this property to construct a linear system and via this linear system, we obtain a upper bound for $g(\hat{\mathbf{p}}_k) - g(\theta_{-k}\mathbf{q}_{\sigma_k})$ in Claim 5.12.

Claim 5.11. *At any equilibrium $s = (s_1, ..., s_n)$, for each agent i, fix s_{-i}, agent i's private signal $\sigma_i \in \Sigma$ and reported signal $\hat{\sigma}_i \in \Sigma$, then there exists a unique prediction which is agent i's best response.*

We define this unique prediction as $\hat{\mathbf{p}}(i, \sigma_i, \hat{\sigma}_i)$

In other words, $s_i(\sigma_i)$ is a distribution over at most m vectors: $\{(\hat{\sigma}_i, \hat{\mathbf{p}}(i, \sigma_i, \hat{\sigma}_i)) | \hat{\sigma}_i \in \Sigma\}$ and

$$Pr_{(\hat{\sigma}_i, \hat{\mathbf{p}}_i) \leftarrow s_i(\sigma)}(\hat{\sigma}_i, \hat{\mathbf{p}}_i) = \begin{cases} \theta_i(\hat{\sigma}_i, \sigma_i) & \hat{\mathbf{p}}_i = \hat{\mathbf{p}}(i, \sigma_i, \hat{\sigma}_i) \\ 0 & \hat{\mathbf{p}}_i \neq \hat{\mathbf{p}}(i, \sigma_i, \hat{\sigma}_i) \end{cases}$$

Proof. For any agent i, assume his private signal is σ_i and he reports $\hat{\sigma}_i$ at equilibrium $(s_1, s_2, ..., s_n)$. Now we will prove there is a unique prediction that maximize agent i's payment.

$$\arg\max_{\hat{\mathbf{p}}} \mathbb{E}[payment(i, \mathcal{M}+)|\sigma_i] \tag{40}$$

$$= \arg\max_{\hat{p}} \alpha PS(\theta_{-i}\mathbf{q}_{\sigma_i}, \hat{\mathbf{p}}) + \beta \sum_{j\neq i} Pr(j) \sum_{\sigma_j} Pr(\sigma_j|\sigma_i) \int_{\hat{\sigma}_j, \hat{\mathbf{p}}_j} Pr_{(\hat{\sigma}_j, \hat{\mathbf{p}}_j)\leftarrow s_j(\sigma_j)}(\hat{\sigma}_j, \hat{\mathbf{p}}_j)\delta(\hat{\sigma}_i = \hat{\sigma}_j) PS(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}) \tag{41}$$

$$= \arg\max_{\hat{p}} \left(\alpha + \beta \sum_{j\neq i} Pr(j) \sum_{\sigma_j} Pr(\sigma_j|\sigma_i) \int_{\hat{\sigma}_j, \hat{\mathbf{p}}_j} Pr_{(\hat{\sigma}_j, \hat{\mathbf{p}}_j)\leftarrow s_j(\sigma_j)}(\hat{\sigma}_j, \hat{\mathbf{p}}_j)\delta(\hat{\sigma}_i = \hat{\sigma}_j) \right)$$
$$PS(\frac{\alpha\theta_{-i}\mathbf{q}_{\sigma_i} + \beta \sum_{j\neq i} Pr(j) \sum_{\sigma_j} Pr(\sigma_j|\sigma_i) \int_{\hat{\sigma}_j, \hat{\mathbf{p}}_j} Pr_{(\hat{\sigma}_j, \hat{\mathbf{p}}_j)\leftarrow s_j(\sigma_j)}(\hat{\sigma}_j, \hat{\mathbf{p}}_j)\delta(\hat{\sigma}_i = \hat{\sigma}_j)\hat{\mathbf{p}}_j}{\alpha + \beta \sum_{j\neq i} Pr(j) \sum_{\sigma_j} Pr(\sigma_j|\sigma_i) \int_{\hat{\sigma}_j, \hat{\mathbf{p}}_j} Pr_{(\hat{\sigma}_j, \hat{\mathbf{p}}_j)\leftarrow s_j(\sigma_j)}(\hat{\sigma}_j, \hat{\mathbf{p}}_j)\delta(\hat{\sigma}_i = \hat{\sigma}_j)}, \hat{\mathbf{p}}) \tag{42}$$

$$= \frac{\alpha\theta_{-i}\mathbf{q}_{\sigma_i} + \beta \sum_{j\neq i} Pr(j) \sum_{\sigma_j} Pr(\sigma_j|\sigma_i) \int_{\hat{\sigma}_j, \hat{\mathbf{p}}_j} Pr_{(\hat{\sigma}_j, \hat{\mathbf{p}}_j)\leftarrow s_j(\sigma_j)}(\hat{\sigma}_j, \hat{\mathbf{p}}_j)\delta(\hat{\sigma}_i = \hat{\sigma}_j)\hat{\mathbf{p}}_j}{\alpha + \beta \sum_{j\neq i} Pr(j) \sum_{\sigma_j} Pr(\sigma_j|\sigma_i) \int_{\hat{\sigma}_j, \hat{\mathbf{p}}_j} Pr_{(\hat{\sigma}_j, \hat{\mathbf{p}}_j)\leftarrow s_j(\sigma_j)}(\hat{\sigma}_j, \hat{\mathbf{p}}_j)\delta(\hat{\sigma}_i = \hat{\sigma}_j)} \tag{43}$$

In equation (41), the first part is the prediction score of agent i, the second part is part of the information score of agent i. Note that for the information score $PS(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}) - PS(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_j)$ of agent i, only $PS(\hat{\mathbf{p}}_j, \hat{\mathbf{p}})$ is related to agent i's reported prediction $\hat{\mathbf{p}}$ so we only consider this part to analyze the equilibrium. $Pr(j)$ is the probability that agent j is matched with agent i, $Pr(\sigma_j|\sigma_i)$ is the probability that agent j receives σ_j given agent i receives σ_i. Then given agent j's strategy s_j and private signal, we integrate over agent j possible report profiles and only consider the case $\hat{\sigma}_i = \hat{\sigma}_j$.

The second equality follows since proper scoring rule is linear for the first entry.

The last equality follows since we obtain the highest value only if $\hat{\mathbf{p}}$ equals the first entry based on the property of strict proper scoring rule.

□

The below claim tells us we can bound the distance between each agent's best response prediction (the prediction which maximizes his total reward) and his best prediction (the prediction which maximizes his prediciton score) by the distance between his best response prediction and his neighbors' best response predictions.

Claim 5.12. *For any convex function $g(\cdot)$, for any σ_i and $\hat{\sigma}_i$, we have*

$$\alpha Pr(\sigma_i)(g(\hat{\mathbf{p}}(i, \sigma_i, \hat{\sigma}_i)) - g(\theta_{-i}\mathbf{q}_{\sigma_i})) \leq \beta \sum_{j\neq i} Pr(j) \sum_{\sigma_j} Pr(\sigma_j, \sigma_i)\theta_j(\hat{\sigma}_i, \sigma_j)(g(\hat{\mathbf{p}}(j, \sigma_j, \hat{\sigma}_i)) - g(\hat{\mathbf{p}}(i, \sigma_i, \hat{\sigma}_i)))$$

Proof. Based on Claim 5.11, we can rewrite (40)=(42) as a $n \times m \times m$ linear system about $\{\hat{\mathbf{p}}(k, \sigma_k, \hat{\sigma}_k)|k \in [1, n], \sigma_k \in \Sigma, \hat{\sigma}_k \in \Sigma\}$:

$$\hat{\mathbf{p}}(i, \sigma_i, \hat{\sigma}_i) = \arg\max_{\hat{\mathbf{p}}} \mathbb{E}[payment(i, \mathcal{M}+)|\sigma_i] \tag{44}$$

$$= \frac{\alpha\theta_{-i}\mathbf{q}_{\sigma_i} + \beta \sum_{j\neq i} Pr(j) \sum_{\sigma_j} Pr(\sigma_j|\sigma_i)\theta_j(\hat{\sigma}_i, \sigma_j)\hat{\mathbf{p}}(j, \sigma_j, \hat{\sigma}_i)}{\alpha + \beta \sum_{j\neq i} Pr(j) \sum_{\sigma_j} Pr(\sigma_j|\sigma_i)\theta_j(\hat{\sigma}_i, \sigma_j)} \tag{45}$$

Fix i, let $\lambda_i = \frac{\alpha}{\alpha + \beta \sum_{j\neq i} Pr(j) \sum_{\sigma_j} Pr(\sigma_j|\sigma_i)\theta_j(\hat{\sigma}_i, \sigma_j)}$, $\lambda_{j,\sigma_j} = \frac{\beta \sum_{j\neq i} Pr(j) \sum_{\sigma_j} Pr(\sigma_j|\sigma_i)\theta_j(\hat{\sigma}_i, \sigma_j)}{\alpha + \beta \sum_{j\neq i} Pr(j) \sum_{\sigma_j} Pr(\sigma_j|\sigma_i)\theta_j(\hat{\sigma}_i, \sigma_j)}$ for $j \neq i$ and $\sigma_j \in \Sigma$, we have $\lambda_i + \sum_{j\neq i, \sigma_j} \lambda_{j,\sigma_j} = 1$

Based on the convexity of $g(\cdot)$, we have

$$\begin{aligned} g(\hat{\mathbf{p}}(i,\sigma_i,\hat{\sigma}_i)) &= g(\lambda_i\theta_{-i}\mathbf{q}_{\sigma_i} + \sum_{j\neq i,\sigma_j}\lambda_{j,\sigma_j}\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_i)) \\ &\leq \lambda_i g(\theta_{-i}\mathbf{q}_{\sigma_i}) + \sum_{j\neq i,\sigma_j}\lambda_{j,\sigma_j} g(\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_i)) \end{aligned}$$

After substitutions, we multiply $\left(\alpha+\beta Pr(j)\sum_{j\neq i}\sum_{\sigma_j}Pr(\sigma_j|\sigma_i)\theta_j(\hat{\sigma}_i,\sigma_j)\right)Pr(\sigma_i)$ in both sides. Note that $Pr(\sigma_i)Pr(\sigma_j|\sigma_i) = Pr(\sigma_j,\sigma_i)$, then by manipulation, the claim follows.

□

Claim 5.12 gives an upper bound to $g(\hat{\mathbf{p}}_k) - g(\theta_{-k}\mathbf{q}_{\sigma_k})$ that is the distance between agent k's best response prediction and his neighbors' best response predictions. Now we continue the proof for our main lemma.

To bound

$$\begin{aligned} &\sum_{j,k,\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k)* \\ &\int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k})(D^*(\hat{\mathbf{p}}_j,\hat{\mathbf{p}}_k) - D^*(\hat{\mathbf{p}}_j,\theta_{-k}\mathbf{q}_{\sigma_k}) + D^*(\hat{\mathbf{p}}_j,\theta_{-k}\mathbf{q}_{\sigma_k}) - D^*(\theta_{-j}\mathbf{q}_{\sigma_j},\theta_{-k}\mathbf{q}_{\sigma_k})) \end{aligned} \tag{46}$$

We rewrite $\int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k})$ as $\theta_j(\hat{\sigma}_j,\sigma_j)\theta_k(\hat{\sigma}_k,\sigma_k)$ and $\hat{\mathbf{p}}_j$ as $\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_j)$, $\hat{\mathbf{p}}_k$ as $\hat{\mathbf{p}}(k,\sigma_k,\hat{\sigma}_k)$ which we can do because of Claim 5.11.

We first give an upper bound to

$$\sum_{j,k}\sum_{\sigma_j,\sigma_k,\hat{\sigma}_j,\hat{\sigma}_k} Pr(j,k)Pr(\sigma_j,\sigma_k)\theta_j(\hat{\sigma}_j,\sigma_j)\theta_k(\hat{\sigma}_k,\sigma_k)(D^*(\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_j),\hat{\mathbf{p}}(k,\sigma_k,\hat{\sigma}_k)) - D^*(\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_j),\theta_{-k}\mathbf{q}_{\sigma_k}))$$

The analysis for the second part is similar.

Based on Claim 5.12, we have

$$\sum_{j,k}\sum_{\sigma_j,\sigma_k,\hat{\sigma}_j,\hat{\sigma}_k} Pr(j,k)Pr(\sigma_j,\sigma_k)\theta_j(\hat{\sigma}_j,\sigma_j)\theta_k(\hat{\sigma}_k,\sigma_k)(D^*(\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_j),\hat{\mathbf{p}}(k,\sigma_k,\hat{\sigma}_k)) - D^*(\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_j),\theta_{-k}\mathbf{q}_{\sigma_k})) \tag{47}$$

$$\leq \sum_{j,k}\sum_{\sigma_j,\sigma_k,\hat{\sigma}_j,\hat{\sigma}_k} Pr(j,k)Pr(\sigma_j,\sigma_k)\theta_j(\hat{\sigma}_j,\sigma_j)\theta_k(\hat{\sigma}_k,\sigma_k) \tag{48}$$

$$\frac{\beta}{\alpha Pr(\sigma_k)}\sum_{l\neq k}\sum_{\sigma_l} Pr(l)Pr(\sigma_l,\sigma_k)\theta_l(\hat{\sigma}_k,\sigma_l)(D^*(\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_j),\hat{\mathbf{p}}(l,\sigma_l,\hat{\sigma}_k)) - D^*(\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_j),\hat{\mathbf{p}}(k,\sigma_k,\hat{\sigma}_k))) \tag{49}$$

Since $\frac{Pr(\sigma_j,\sigma_k)}{Pr(\sigma_k)} \leq 1$, we obtain (50) from (49).

$$(49) \leq \sum_{j,k} \sum_{\sigma_j,\sigma_k,\hat{\sigma}_j,\hat{\sigma}_k} Pr(j,k)\theta_j(\hat{\sigma}_j,\sigma_j)\theta_k(\hat{\sigma}_k,\sigma_k)$$
$$\frac{\beta}{\alpha}\sum_{l\neq k}\sum_{\sigma_l} Pr(l)Pr(\sigma_l,\sigma_k)\theta_l(\hat{\sigma}_k,\sigma_l)(D^*(\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_j),\hat{\mathbf{p}}(l,\sigma_l,\hat{\sigma}_k)) - D^*(\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_j),\hat{\mathbf{p}}(k,\sigma_k,\hat{\sigma}_k))) \tag{50}$$

$$\leq \sum_{j,k} \sum_{\sigma_j,\sigma_k,\hat{\sigma}_j,\hat{\sigma}_k} Pr(j,k)\theta_j(\hat{\sigma}_j,\sigma_j)\theta_k(\hat{\sigma}_k,\sigma_k)$$
$$\frac{\beta}{\alpha}\sum_{l\neq k}\sum_{\sigma_l} Pr(l)Pr(\sigma_l,\sigma_k)\theta_l(\hat{\sigma}_k,\sigma_l)|(D^*(\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_j),\hat{\mathbf{p}}(l,\sigma_l,\hat{\sigma}_k)) - D^*(\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_j),\hat{\mathbf{p}}(k,\sigma_k,\hat{\sigma}_k)))| \tag{51}$$

Note that (51) and (50) are identical except for the value sign.

Then we obtain (52) from (51) since

$$\begin{aligned}|D^*(x,y) - D^*(x,z)| &\leq (\sqrt{D^*(x,y)} + \sqrt{D^*(x,z)})|\sqrt{D^*(x,y)} - \sqrt{D^*(x,z)})| \\ &\leq 2|\sqrt{D^*(x,y)} - \sqrt{D^*(x,z)})|\end{aligned}$$

The second inequality follows since $0 \leq D^* \leq 1$

$$(51) \leq 2\sum_{j,k} \sum_{\sigma_j,\sigma_k,\hat{\sigma}_j,\hat{\sigma}_k} Pr(j,k)\theta_j(\hat{\sigma}_j,\sigma_j)\theta_k(\hat{\sigma}_k,\sigma_k)$$
$$\frac{\beta}{\alpha}\sum_{l\neq k}\sum_{\sigma_l} Pr(l)Pr(\sigma_l,\sigma_k)\theta_l(\hat{\sigma}_k,\sigma_l)$$
$$|(\sqrt{D^*(\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_j),\hat{\mathbf{p}}(l,\sigma_l,\hat{\sigma}_k))} - \sqrt{D^*(\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_j),\hat{\mathbf{p}}(k,\sigma_k,\hat{\sigma}_k)))}| \tag{52}$$

Once we get (52), we can use the fact that $\sqrt{D^*}$ is metric which implies the triangle inequality; (53) follows.

$$(52) \leq 2\sum_{j,k} \sum_{\sigma_j,\sigma_k,\hat{\sigma}_j,\hat{\sigma}_k} Pr(j,k)\theta_j(\hat{\sigma}_j,\sigma_j)\theta_k(\hat{\sigma}_k,\sigma_k)$$
$$\frac{\beta}{\alpha}\sum_{l\neq k}\sum_{\sigma_l} Pr(l)Pr(\sigma_l,\sigma_k)\theta_l(\hat{\sigma}_k,\sigma_l)(\sqrt{D^*(\hat{\mathbf{p}}(k,\sigma_k,\hat{\sigma}_k),\hat{\mathbf{p}}(l,\sigma_l,\hat{\sigma}_k))}) \tag{53}$$

Note that $\sum_{\sigma_j}\sum_{\hat{\sigma}_j}\theta_j(\hat{\sigma}_j,\sigma_j) = \sum_{\sigma_j} 1 = m$, also we have $\sum_j Pr(l) = \sum_j Pr(j) = 1$, $Pr(k,l) = Pr(j,k)$ then (54) follows.

$$(53) = 2m\frac{\beta}{\alpha}\sum_{l}\sum_{k\neq l}\sum_{\sigma_k,\hat{\sigma}_k} Pr(k,l)\theta_k(\hat{\sigma}_k,\sigma_k)\sum_{\sigma_l} Pr(\sigma_l,\sigma_k)\theta_l(\hat{\sigma}_k,\sigma_l)(\sqrt{D^*(\hat{\mathbf{p}}(k,\sigma_k,\hat{\sigma}_k),\hat{\mathbf{p}}(l,\sigma_l,\hat{\sigma}_k))}) \tag{54}$$

$$= 2m\frac{\beta}{\alpha}\sum_{k,l\neq k}\sum_{\sigma_k,\hat{\sigma}_k,\sigma_l} Pr(k,l)\theta_k(\hat{\sigma}_k,\sigma_k)Pr(\sigma_l,\sigma_k)\theta_l(\hat{\sigma}_k,\sigma_l)(\sqrt{D^*(\hat{\mathbf{p}}(k,\sigma_k,\hat{\sigma}_k),\hat{\mathbf{p}}(l,\sigma_l,\hat{\sigma}_k))}) \tag{55}$$

$$= 2m\frac{\beta}{\alpha} \times Inconsistency \tag{56}$$

The analysis for the second part

$$\sum_{j,k}\sum_{\sigma_j,\sigma_k,\hat{\sigma}_j,\hat{\sigma}_k} Pr(j,k)Pr(\sigma_j,\sigma_k)\theta_j(\hat{\sigma}_j,\sigma_j)\theta_k(\hat{\sigma}_k,\sigma_k)(D^*(\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_j),\theta_{-k}\mathbf{q}_{\sigma_k}) - D^*(\theta_{-j}\mathbf{q}_{\sigma_j},\theta_{-k}\mathbf{q}_{\sigma_k}))$$

is similar, note that j and k are symmetric and $D^*(\cdot,\theta_{-k}\mathbf{q}_{\sigma_k})$ is a convex function. We can use Claim 5.12 and triangle inequality to bound the second part by $2m\frac{\beta}{\alpha}\times$ *Inconsistency*.

So if we set $2m\frac{\beta}{\alpha}<\frac{1}{2}$, then *TotalDivergence*$(s)-$*TotalDivergence*$(s_{BP})<$ *Inconsistency*, proving the inequality in our main lemma.

To prove that if the equality in our main lemma holds then $s=s_{BP}$, we first show that

Claim 5.13. *The equality in ClassificationScore*$(s)\leq$ *TotalDivergence*(s_{BP}) *holds iff Inconsistency*$(s)=0$.

Proof. Note that (36) tells us when *ClassificationScore*$(s)=$ *TotalDivergence*(s_{BP}), we have *Diversity*$(s)=$ *TotalDivergence*(s) which implies *Inconsistency*$(s)=0$ based on Claim 5.4. □

Then we will prove

Claim 5.14. *If Inconsistency*$(s)=0$ *then* $s=s_{BP}$

Proof. Recall in (44), we have for any i,

$$\begin{aligned}\hat{\mathbf{p}}(i,\sigma_i,\hat{\sigma}_i) &= \arg\max_{\hat{\mathbf{p}}}\mathbb{E}[payment(i,\mathcal{M}+)|\sigma_i]\\ &= \frac{\alpha\theta_{-i}\mathbf{q}_{\sigma_i}+\beta Pr(j)\sum_{j\neq i}\sum_{\sigma_j}Pr(\sigma_j|\sigma_i)\theta_j(\hat{\sigma}_i,\sigma_j)\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_i)}{\alpha+\beta Pr(j)\sum_{j\neq i}\sum_{\sigma_j}Pr(\sigma_j|\sigma_i)\theta_j(\hat{\sigma}_i,\sigma_j)}\end{aligned} \tag{57}$$

If *Inconsistency*$(s)=0$, we can see if $\theta_j(\hat{\sigma}_i,\sigma_j)>0$ we must have $\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_i)=\hat{\mathbf{p}}_i(i,\sigma_i,\hat{\sigma}_i)$. So we have $\theta_{-i}\mathbf{q}_{\sigma_i}=\hat{\mathbf{p}}_i(i,\sigma_i,\hat{\sigma}_i)$ for any i since we have $\alpha\theta_{-i}\mathbf{q}_{\sigma_i}=\alpha\hat{\mathbf{p}}_i(i,\sigma_i,\hat{\sigma}_i)$ if we multiply $\alpha+\beta Pr(j)\sum_{j\neq i}\sum_{\sigma_j}Pr(\sigma_j|\sigma_i)\theta_j(\hat{\sigma}_i,\sigma_j)$ in both sides of equation (57) and combine the fact that $\hat{\mathbf{p}}(j,\sigma_j,\hat{\sigma}_i)=\hat{\mathbf{p}}_i(i,\sigma_i,\hat{\sigma}_i)$. Thus we have $s=s_{BP}$. □

5.4 Proof for Claims

Claim 2.10. *Assume that the distribution over all agents' private signals is* $\omega\in\Delta_\Sigma$*, the distribution over all agents' reported signals will be* $\bar{\theta}_n\omega$.

Proof for Claim 2.10. The probability of signal σ will be

$$\sum_i Pr(i)\sum_{\sigma'}\theta_i(\sigma,\sigma')\omega(\sigma') = \frac{1}{n}\sum_i\sum_{\sigma'}\theta_i(\sigma,\sigma')\omega(\sigma') = \sum_{\sigma'}\bar{\theta}_n(\sigma,\sigma')\omega(\sigma')$$

where $Pr(i)$ is the probability agent i is picked. For each agent i, we sum the probability agent i receives private signal σ' which is $\omega(\sigma')$ times the probability that he reports σ given he receives σ' which is $\theta_i(\sigma,\sigma')$ over all possible private signal σ'.

So the distribution of reported signals is $\bar{\theta}_n\omega$. □

Claim 2.12. *For each agent* i*, if he receives private signal* σ_i*, agent* i *will believe that the expected likelihood of other agents' reported signals is* $\theta_{-i}\mathbf{q}_{\sigma_i}$ *where* $\theta_{-i}=\frac{\sum_{j\neq i}\theta_j}{n-1}$.

Proof for Claim 2.12. For each agent i, given he receives private signal σ_i, he will believe the expected likelihood for other agents' private signals is $\mathbf{q}_{\sigma_i}$. Based on Claim 2.10, he will believe the expected likelihood for other agents' reported signals is the average signal strategy of other agents' signal strategies times $\mathbf{q}_{\sigma_i}$ which is $\theta_{-i}\mathbf{q}_{\sigma_i}$ where $\theta_{-i} = \frac{\sum_{j\neq i}\theta_j}{n-1}$. □

Claim 2.16. *For any transition matrix $\theta_{m\times m}$ where the sum of every column is 1, θ is a permutation matrix iff for any row of θ, there at most one non-zero entry.*

Proof for Claim 2.16. It is clear that any permutation matrix has exactly one non-zero entry, which is 1, in each row and each column. Thus we only need to prove the direction that if for any row of θ, there is at most one non-zero entry, θ must be a permutation matrix.

We first prove that there are exactly m non-zero entries in θ: if for any row of θ, there is at most one non-zero entry, we can see θ has at most m non-zero entries. θ is a transition matrix where the sum of every column is 1,which implies that θ has at least m non-zero entries. Thus we proved there are exactly m non-zero entries in θ.

We have just shown that θ has exactly m non-zero entries. Since θ has at most one non-zero entry in each row, θ must have exactly one non-zero entry in each row. θ also has at least one non-zero entry in each column since it is a transition matrix, so θ must have exactly one non-zero entry 1 in each column. Thus θ has exactly one non-zero entry 1 in each row and each column which implies that θ is a permutation matrix. □

Claim 4.4. *The* Disagreement Mechanism *has the same equilibria as the* Truthful Mechanism*.*

Proof for Claim 4.4. The value of $score_C(r_j, r_k)$ does not depend on agent i's strategy. The term related to agent i's strategy contained in $score_{\mathcal{M}}$ is $payment_{\mathcal{M}(\alpha,\beta,PS(\cdot,\cdot))}(i,\mathbf{r})$. This implies that agent i's marginal benefit from deviation in $\mathcal{M}+(\alpha,\beta,PS(\cdot,\cdot))$ is the same with its marginal benefit from the same deviation in $\mathcal{M}(\alpha,\beta,PS(\cdot,\cdot))$. □

Claim 5.1.

$$ClassificationScore = Diversity - Inconsistency$$

Proof for Claim 5.1. Based on the definition of *ClassificationScore*, we have

$$\sum_{\substack{i\\ j\neq i}} \sum_{\substack{k\neq i,j\\ \sigma_i,\sigma_j,\sigma_k}} Pr(i)Pr(\sigma_i)Pr(j,k)Pr(\sigma_j,\sigma_k|\sigma_i)* \tag{58}$$

$$\int_{\hat{\sigma}_j,\hat{\mathbf{p}}_j,\hat{\sigma}_k,\hat{\mathbf{p}}_k} Pr_{(\hat{\sigma}_j,\hat{\mathbf{p}}_j)\leftarrow s_j(\sigma_j)}(\hat{\sigma}_j,\hat{\mathbf{p}}_j)Pr_{(\hat{\sigma}_k,\hat{\mathbf{p}}_k)\leftarrow s_k(\sigma_k)}(\hat{\sigma}_k,\hat{\mathbf{p}}_k)score_C(r_j,r_k) \tag{59}$$

Now we begin our proof:

$$
\begin{aligned}
&\sum_{\substack{i\\ j\neq i}}\sum_{\substack{k\neq i,j\\ \sigma_i,\sigma_j,\sigma_k}} Pr(i)Pr(\sigma_i)Pr(j,k)Pr(\sigma_j,\sigma_k|\sigma_i)\int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k})score_C(r_j,r_k)\\
=&\sum_{\substack{i\\ j\neq i}}\sum_{\substack{k\neq i,j\\ \sigma_j,\sigma_k}} \frac{1}{n(n-1)(n-2)}Pr(\sigma_j,\sigma_k)\int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k})score_C(r_j,r_k)\\
=&\frac{1}{n(n-1)}\sum_{\substack{j\\ k\neq j}}\sum_{\sigma_j,\sigma_k} Pr(\sigma_j,\sigma_k)\int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k})score_C(r_j,r_k)\\
=&\sum_{\substack{j\\ k\neq j}}\sum_{\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k)\int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k})score_C(r_j,r_k)
\end{aligned}
$$

The first equality follows since fix j,k, $score_C(r_j,r_k)$ does not depend on i and we also have $\sum_{\sigma_i} Pr(\sigma_i)Pr(\sigma_j,\sigma_k|\sigma_i) = Pr(\sigma_j,\sigma_k)$.

The second equality follows since for any $(j,k), j\neq k$ pair, there are $n-2$ numbers that are neither j nor k which means (j,k) will repeat $n-2$ times since there are $n-2$ possible i.

By definition we can see $\textit{ClassificationScore} = \textit{Diversity} - \textit{Inconsistency}$. □

Claim 5.2. *Any permutation strategy profile has the same ClassificationScore, Diversity, and Inconsistency with truth-telling.*

Proof for Claim 5.2. Any permutation strategy profile's report profiles can be seen as a relabeling to truth-telling's report profiles, which implies the claim. □

Claim 5.3. *The average agent-welfare in our* Disagreement Mechanism *is ClassificationScore*

Proof for Claim 5.3. We only need to prove $\sum_i score_{\mathcal{M}}(i,\mathbf{r}) = 0$.

$$
\begin{aligned}
\sum_i score_{\mathcal{M}}(i,\mathbf{r}) =& \sum_{i\in A} score_{\mathcal{M}}(i,\mathbf{r}) + \sum_{i\in B} score_{\mathcal{M}}(i,\mathbf{r})\\
=&\sum_{i\in A}\left(payment_{\mathcal{M}(\alpha,\beta,PS(\cdot,\cdot))}(i,\mathbf{r}) - \frac{1}{|A|}\sum_{i\in B} payment_{\mathcal{M}(\alpha,\beta,PS(\cdot,\cdot))}(i,\mathbf{r})\right)\\
&+\sum_{i\in B}\left(payment_{\mathcal{M}(\alpha,\beta,PS(\cdot,\cdot))}(i,\mathbf{r}) - \frac{1}{|B|}\sum_{i\in A} payment_{\mathcal{M}(\alpha,\beta,PS(\cdot,\cdot))}(i,\mathbf{r})\right) = 0
\end{aligned}
$$

□

Claim 5.4. *For any strategy profile s, $\textit{Diversity}(s) = \textit{TotalDivergence}(s) \Leftrightarrow \textit{Inconsistency}(s) = 0$*

Proof for Claim 5.4. Note that

$$
\begin{aligned}
&\textit{TotalDivergence}(s) - \textit{Diversity}(s)\\
=&\sum_{\substack{j\\ k\neq j}}\sum_{\sigma_j,\sigma_k} Pr(j,k)Pr(\sigma_j,\sigma_k)\int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k})\delta(\hat{\sigma}_j=\hat{\sigma}_k)D^*(\hat{\mathbf{p}}_j,\hat{\mathbf{p}}_k)
\end{aligned}
$$

while

$$Inconsistency(s) = \sum_{\substack{j \\ k \neq j}} \sum_{\sigma_j, \sigma_k} Pr(j,k) Pr(\sigma_j, \sigma_k) \int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k}) \delta(\hat{\sigma}_j = \hat{\sigma}_k) \sqrt{D^*(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_k)}$$

Because each part in $TotalDivergence(s) - Diversity(s)$ is non-negative, $TotalDivergence(s) - Diversity(s) = 0$ will imply $Pr(\sigma_j, \sigma_k) \int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k}) \delta(\hat{\sigma}_j = \hat{\sigma}_k) D^*(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_k) = 0$.

So we have $Pr(\hat{j},\hat{k}) \delta(\hat{\sigma}_j = \hat{\sigma}_k) = 0$ or $D^*(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_k) = 0$ which implies $Pr(\sigma_j, \sigma_k) \int_{\hat{j},\hat{k}} Pr(\hat{j},\hat{k}) \delta(\hat{\sigma}_j = \hat{\sigma}_k) \sqrt{D^*(\hat{\mathbf{p}}_j, \hat{\mathbf{p}}_k)} = 0$. The proof for another direction is similar. □

6 Conclusion

We have shown that our *Disagreement Mechanism* promotes truth-telling by 1) having truth-telling as a Bayesian Nash equilibrium; 2) having no other symmetric equilibrium with agent welfare more than the truth-telling equilibrium; 3) having the agent welfare of any equilibrium approach that of truth-telling as the number of agents increases; 4) requiring that any equilibrium with agent welfare close to that of truth-telling must be close to a permutation strategy.

We have argued, that our mechanism is near optimal in the sense that no truthful mechanism without knowledge of the common prior can avoid having permutation equilibrium with high agent welfare, and, in our mechanism, any equilibrium with agent welfare even close to that of truth-telling must be close to a permutation equilibrium.

Permutation equilibria are intuitively unnatural and risky as they require extreme coordination amongst the agents. We believe they are very unlikely to occur in practice. Additionally, any asymmetric equilibrium also seem unlikely, especially as the number of agents increases because a) such deviations help less as the number of agents' increases; b) implementing such deviations will become increasingly difficult as the number of agents increases. Thus our results about symmetric equilibria and equilibrium in general are quite strong, despite the impossibility result.

In addition to the above results, our work has several contributions in the techniques employed:

(1) In a common prior setting, agents with the same private information cannot agree to disagree. Thus, agents with the same signal should report similarly. Our *Disagreement Mechanism* encourages not only agents with the same private information to agree, but also agents with different private information to disagree. We do this by employing tools from information theory, namely *Information Monotonicity*. Despite their natural and powerful application, to our knowledge, this is the first time such tools have been explicitly employed in the peer prediction literature.

(2) Additionally, we created a framework for understanding the space of equilibrium which was integral to our results.

We hope that both the information theory tools and the new understanding of equilibrium introduced in this work will continue to provide useful insights for designing and analyzing future mechanisms in peer-prediction and related settings.

References

[1] S-I Amari and A Cichocki. Information geometry of divergence functions. *Bulletin of the Polish Academy of Sciences: Technical Sciences*, 58(1):183–195, 2010.

[2] Shun-ichi Amari and Hiroshi Nagaoka. *Methods of information geometry*, volume 191. American Mathematical Soc., 2007.

[3] Yang Cai, Constantinos Daskalakis, and Christos H Papadimitriou. Optimum statistical estimation with strategic data sources. *arXiv preprint arXiv:1408.2539*, 2014.

[4] Anirban Dasgupta and Arpita Ghosh. Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd international conference on World Wide Web*, pages 319–330. International World Wide Web Conferences Steering Committee, 2013.

[5] Boi Faltings, Radu Jurca, Pearl Pu, and Bao Duy Tran. Incentives to counter bias in human computation. In *Second AAAI Conference on Human Computation and Crowdsourcing*, 2014.

[6] Xi Alice Gao, Andrew Mao, Yiling Chen, and Ryan Prescott Adams. Trick or treat: putting peer prediction to the test. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 507–524. ACM, 2014.

[7] Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.

[8] Radu Jurca and Boi Faltings. Collusion-resistant, incentive-compatible feedback payments. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 200–209. ACM, 2007.

[9] Radu Jurca and Boi Faltings. Mechanisms for making crowds truthful. *J. Artif. Int. Res.*, 34 (1), March 2009.

[10] Vijay Kamble, Nihar Shah, David Marn, Abhay Parekh, and Kannan Ramachandran. Truth serums for massively crowdsourced evaluation tasks. *arXiv preprint arXiv:1507.07045*, 2015.

[11] Y. Kong, G. Schoenebeck, and K. Ligett. Putting Peer Prediction Under the Micro(economic)scope and Making Truth-telling Focal. *ArXiv e-prints*, March 2016.

[12] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, pages 1359–1373, 2005.

[13] Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013.

[14] D. Prelec. A Bayesian Truth Serum for subjective data. *Science*, 306(5695):462–466, 2004.

[15] Goran Radanovic and Boi Faltings. A robust bayesian truth serum for non-binary signals. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence, AAAI 2013*, number EPFL-CONF-197486, pages 833–839, 2013.

[16] Goran Radanovic and Boi Faltings. Incentives for truthful information elicitation of continuous signals. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.

[17] Blake Riley. Minimum truth serums with optional predictions. In *Proceedings of the 4th Workshop on Social Computing and User Generated Content (SC14)*, 2014.

[18] J. Witkowski and D. Parkes. A robust Bayesian Truth Serum for small populations. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI 2012)*.

[19] Jens Witkowski and David C Parkes. Peer prediction without a common prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 964–981. ACM, 2012.

[20] Jens Witkowski and David C Parkes. Learning the prior in minimal peer prediction. In *Proceedings of the 3rd Workshop on Social Computing and User Generated Content at the ACM Conference on Electronic Commerce*, page 14. Citeseer, 2013.

[21] Jens Witkowski, Bernhard Nebel, and David C Parkes. *Robust Peer Prediction Mechanisms*. PhD thesis, Ph. D. Dissertation, Albert-Ludwigs-Universitat Freiburg: Institut fur Informatik, 2014.

[22] Peter Zhang and Yiling Chen. Elicitability and knowledge-free elicitation with peer prediction. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 245–252. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

7 Appendix

To understand the strictness condition more in Lemma 2.18, we give an example where the strictness condition is not satisfied:

Example 7.1. $\mathbf{p} = (0.1\ 0.2\ 0.7)$, $\mathbf{q} = (0.2\ 0.4\ 0.4)$, $\theta = \begin{pmatrix} 0.3 & 0.6 & 0 \\ 0.7 & 0.4 & 0 \\ 0 & 0 & 1 \end{pmatrix}$.

We show by case analysis that we cannot find $\sigma, \sigma', \sigma''$ *such that* $\theta(\sigma, \sigma')\mathbf{p}(\sigma') > 0$, $\theta(\sigma, \sigma'')\mathbf{p}(\sigma'') > 0$ *and* $\frac{\mathbf{p}(\sigma'')}{\mathbf{p}(\sigma')} \neq \frac{\mathbf{q}(\sigma'')}{\mathbf{q}(\sigma')}$.

First note that because $\frac{\mathbf{p}(\sigma'')}{\mathbf{p}(\sigma')} \neq \frac{\mathbf{q}(\sigma'')}{\mathbf{q}(\sigma')}$, *it cannot be that* $\sigma' = \sigma''$, *nor can it be the case that* $\sigma', \sigma'' \in \{1, 2\}$ *because* $\frac{\mathbf{p}(1)}{\mathbf{p}(2)} = \frac{\mathbf{q}(1)}{\mathbf{q}(2)}$ *and* $\frac{\mathbf{p}(2)}{\mathbf{p}(1)} = \frac{\mathbf{q}(2)}{\mathbf{q}(1)}$. *Thus it must be that either* $\sigma' \in \{1, 2\}$ *and* $\sigma'' = 3$ *or* $\sigma' = 3$ *and* $\sigma'' \in \{1, 2\}$. *Because these are symmetric, we consider the first case.*

Because $\theta(\sigma, \sigma')\mathbf{p}(\sigma') > 0$ *it must be that* $\sigma \in \{1, 2\}$, *but because* $\theta(\sigma, \sigma'')\mathbf{p}(\sigma'') > 0$, *it must be that* $\sigma = 3$. *So no assignment of* $\sigma, \sigma', \sigma''$ *is possible.*

Thus, the strictness condition is not satisfied. By simple calculations, we have $\theta\mathbf{p} = (0.15\ 0.15\ 0.7)$, $\theta\mathbf{q} = (0.3\ 0.3\ 0.4)$. *By some algebraic calculations, we have* $D_f(\mathbf{p}, \mathbf{q}) = D_f(\theta\mathbf{p}, \theta\mathbf{q})$ *for any function* f.